

Junos® OS

Ethernet Interfaces User Guide for Routing Devices

Published
2026-01-23

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

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

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

Junos® OS Ethernet Interfaces User Guide for Routing Devices

Copyright © 2026 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About This Guide | **xiii**

1

Ethernet Interfaces

Ethernet Interfaces Overview | **2**

Configure Ethernet Interfaces | **4**

Configure Ethernet Interfaces Properties | **4**

Flow Control for Ethernet Interfaces | **6**

Flow Control Overview | **6**

MAC Address Filtering and Accounting on Ethernet Interfaces | **7**

MAC Address Filtering Overview | **8**

Configure MAC Address Filtering for Ethernet Interfaces | **8**

MAC Address Accounting for Ethernet Interfaces | **9**

Management Ethernet Interfaces | **9**

Management Ethernet Interface Overview | **10**

Configure a Consistent Management IP Address | **11**

Configure the MAC Address on the Management Ethernet Interface | **12**

Platform-Specific Management Ethernet Interfaces Behavior | **12**

Point-to-Point Protocol over Ethernet (PPPoE) | **13**

PPPoE Overview | **14**

2

Aggregated Ethernet Interfaces

Aggregated Ethernet Interfaces Overview | **19**

Configure Aggregated Ethernet Interfaces | **20**

Enhanced Link Aggregation Group | **22**

Mixed-Rate Aggregated Ethernet Interfaces | **23**

Link Aggregation Control Protocol | **25**

Configure LACP | 26

Targeted Distribution Across Aggregated Ethernet Member Links | 28

Configure Targeted Distribution Across Aggregated Ethernet Member Links | 29

 Requirements | 33

 Overview | 33

 Verification | 34

MAC Address on Aggregated Ethernet Interfaces | 40

Link Protection of Aggregated Ethernet Interfaces | 42

Aggregated Ethernet Link Protection | 42

 Configure Link Protection for Aggregated Ethernet Interfaces | 43

 Configure Primary and Backup Links for Link Aggregated Ethernet Interfaces | 43

 Revert Traffic on Aggregate Ethernet Interfaces to a Primary Link When Traffic is Passing
 Through a Backup Link | 43

 Disable Link Protection for Aggregated Ethernet Interfaces | 44

Configure Aggregated Ethernet Minimum Links | 44

Example: Configure Aggregated Ethernet Link Protection | 45

Scheduling on Aggregated Ethernet Interfaces | 46

Configure Shared Scheduling on Aggregated Ethernet Interfaces | 46

Configure Scheduler on Aggregated Ethernet Interfaces Without Link Protection | 47

Load Balancing on Aggregated Ethernet Interfaces | 48

Load Balancing and Ethernet Link Aggregation Overview | 49

Understanding Aggregated Ethernet Load Balancing | 49

Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data | 51

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | 54

Configuring Adaptive Load Balancing | 56

Understanding Symmetric Hashing for Load Balancing | 57

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series
Routers | 58

 Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview | 58

Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers | 59

Configuring Symmetrical Load Balancing on Trio-Based MPCs | 62

Example Configurations | 64

Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers | 66

Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers | 68

Example: Configuring Aggregated Ethernet Load Balancing | 71

Example: Configuring Aggregated Ethernet Load Balancing | 71

Requirements | 72

Overview | 72

Configuration | 74

Verification | 88

Platform-Specific Aggregated Ethernet Load Balancing Behavior | 89

Performance Monitoring on Aggregated Ethernet Interfaces | 91

ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces | 91

Guidelines to Configure Performance Monitoring on Aggregated Ethernet Interfaces | 92

Periodic Packet Management on Ethernet Interfaces | 93

Periodic Packet Management Overview | 93

Configure Periodic Packet Management | 94

Enable Centralized Periodic Packet Management | 94

Ethernet Link Aggregation | 95

3

Gigabit Ethernet Interfaces

Gigabit Ethernet PICs | 103

P2-10G-40G-QSFPP PIC Overview | 103

Configure the P2-10G-40G-QSFPP PIC | 114

Configure the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 114

Configure the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 115

Configure the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 115

Configure the PIC at Port Group Level | 116

Configure Framing Mode on P2-10G-40G-QSFPP PIC | 117

Configure LAN PHY or WAN PHY Framing Mode in 10-Gigabit Ethernet Mode | 117

Configure LAN PHY Framing Mode in 40-Gigabit Ethernet Mode | 118

Example: Configure the P2-10G-40G-QSFPP PIC | 118

Requirements | 118

Overview | 119

Configuration | 119

Verification | 121

Framing Overview | 122

Understanding WAN Framing | 123

Configure Ethernet Framing | 123

Gigabit Ethernet MICs | 124

100-Gigabit Ethernet Interfaces Interoperability | 125

Platform-Specific 100-Gigabit Ethernet Interfaces Behavior | 125

Smart SFPs To Transport Legacy Network Traffic over Packet Switched Networks | 127

Smart SFPs To Transport Legacy Traffic | 127

Smart-SFPs for transporting legacy PDH Traffic | 130

Requirements | 130

Overview | 130

Configure the DS3 Smart SFP | 130

Verification | 133

Configure the Smart-SFPs to transport legacy SDH Traffic | 136

Requirements | 136

Overview | 136

Configure the STM1 Smart SFP | 137

Verification | 139

Layer 2 Overhead Attribute in Interface Statistics | 142

Accounting of the Layer 2 Overhead Attribute in Interface Statistics | 142

Configure Layer 2 Overhead Accounting in Interface Statistics | 145

Enable the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level | 146

Verify the Accounting of Layer 2 Overhead in Interface Statistics | 147

Gigabit Ethernet Policers | 150

Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 151

Configure Gigabit Ethernet Policers | 153

 Overview | 153

 Configure a Policer | 154

 Specify an Input Priority Map | 155

 Specify an Output Priority Map | 155

 Apply a Policer | 156

 Example: Configure Gigabit Ethernet Policers | 157

 Example | 157

 Example Configuration | 158

Configure Gigabit Ethernet Two-Color and Tricolor Policers | 159

 Overview | 160

 Configure a Policer | 161

 Apply a Policer | 162

 Configure and Apply a Policer | 162

Gigabit Ethernet Autonegotiation | 164

Gigabit Ethernet Autonegotiation Overview | 164

Configure Gigabit Ethernet Autonegotiation | 165

 Configure Gigabit Ethernet Autonegotiation with Remote Fault | 165

 Configure Autonegotiation Speed | 165

 Autonegotiation Status Display | 166

4

Port Speed

Port Speed Overview | 173

Introduction to Port Speed | 173

Platform-Specific Port Profile Behavior | 175

Configure Port Speed | 176

Port Speed on ACX Routers | 182

Port Speed on ACX710 Router | 182

Port Speed on ACX5448 Router | 184

Port Speed on ACX7020 187
Port Speed on ACX7024 and ACX7024X Router 188
Port Speed on ACX7100 Router 190
Port Speed on ACX7332 and ACX7348 Routers 197
Port Speed on ACX7509 Router 202
Port Speed on MX Routers 210
Port Speed on MX204 211
Port speed on MX301 Router 216
Port speed on MX304 Router 221
Port Speed on MPC7E-MRATE 224
Port Speed on MPC7E-10G 227
Port Speed on MIC-MRATE 229
Port Speed on MIC-3D-10GE-SFP-E 233
Port Speed on MX10003 MPC 234
Port Speed on MX10K-LC4800 240
Port Speed on MIC-MACSEC-20GE 244
Port Speed on MPC10E-10C-MRATE 245
Port Speed on MPC10E-15C-MRATE 250
Port Speed on MX2K-MPC11E 256
Port Speed on MX10K-LC480 266
Port Speed on MX10K-LC2101 268
Port Speed on MX10K-LC9600 272
Port Speed on PTX Routers 274
Port Speed on PTX10K-LC1201 275
Port Speed on PTX10001-36MR Router 283
Port Speed on PTX10K-LC1202-36MR 293

Port Speed on PTX10003 | 301

5

800ZR and 800G OpenZR+ Optical Transceivers

Overview | 309

Features of 800ZR and 800G OpenZR+ | 310

Application Selection for 800G OpenZR+ | 311

Configure Application Selection on 800G OpenZR+ Optics | 315

Transmit Optical Power for 800G OpenZR+ | 319

Configure Transmit Output Power | 320

Optical Loopback | 322

Signal Power | 323

Configure and Monitor 800ZR and 800G OpenZR+ Coherent Optics | 323

6

400ZR and 400G OpenZR+ Optical Transceivers

Overview | 327

Features of 400ZR and 400G OpenZR+ | 329

Application Selection | 329

Configure Application Selection | 332

Target Output Power | 335

Configure Target Output Power | 336

Enhanced Loopback Options | 337

Tunable Wavelength | 338

Signal Power Performance | 338

Ports in High-Power Mode | 341

Configure 400ZR and 400G OpenZR+ Coherent Optics | 342

Coherent Optics Performance Monitoring | 343

User Configurable PM Interval Length | 345

Configure User Configurable PM Interval Length | 346

Configure and Monitor 400ZR and 400G OpenZR+ Coherent Optics | 347

7

100GbE, 40GbE, and 10GbE Optical Transceivers

100GbE, 40GbE, and 10GbE | 351

Overview | 351

Supported Features | 353

Supported PICs | 357

Configure Optics | 358

Supported OTN and Optics Options | 363

Supported OTN Options on ACX6360 and ACX5448 Routers | 364

Supported OTN Options on MX Series Routers | 369

Supported OTN Options on PTX Series Routers | 379

Supported Optics Options on ACX6360 and ACX5448-D Routers | 391

Supported Optics Options on PTX10008 and PTX10016 Series Routers | 395

Forward Error Correction and Bit Error Rate | 400

Overview | 401

Supported Forward Error Correction Modes | 405

ODU Path Delay Measurement for Performance Monitoring | 407

Overview | 407

Enabling ODU Path Delay Measurement | 408

Disabling ODU Path Delay Measurement | 410

| 412

| 412

| 412

| 412

| 412

Interface Mapping and Modulation format for PTX10K-LC1104 Line Card | 413

8

Monitoring and Troubleshooting Ethernet Interfaces

Monitor Interfaces | 417

Passive Monitoring on Ethernet Interfaces | 417

Enable Passive Monitoring on Ethernet Interfaces | 419

Link Degrade Monitoring | 421

Monitor Gigabit Ethernet Interfaces | 422

Checklist to Monitor Gigabit Ethernet Interfaces | 423

How to Monitor Gigabit Ethernet Interfaces | 424

Display the Status of Gigabit Ethernet Interfaces | 424

Display the Status of Gigabit Ethernet Interfaces | 426

Display the Status of a Specific Gigabit Ethernet Interface | 428

Display Extensive Status for a Specific Gigabit Ethernet Interface | 429

Monitor Statistics for a Gigabit Ethernet Interface | 434

Fiber-Optic Ethernet Interface Specifications | 435

Perform Loopback Test for Gigabit Ethernet Interfaces | 437

Checklist to Use Loopback Test for Gigabit Ethernet Interfaces | 437

Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | 439

Create a Loopback | 439

Create a Physical Loopback for a Fiber-Optic Interface | 440

Create a Loopback Plug for an RJ-45 Ethernet Interface | 441

Configure a Local Loopback | 442

Verify That the Gigabit Ethernet Interface Is Up | 444

Configure a Static Address Resolution Protocol Table Entry | 448

Clear Gigabit Ethernet Interface Statistics | 454

Ping the Gigabit Ethernet Interface | 454

Check for Gigabit Ethernet Interface Error Statistics | 456

Diagnose a Suspected Circuit Problem | 459

Troubleshoot Interfaces | 461

Configure Interface Diagnostics Tools to Test the Physical Layer Connections | 461

Configure Loopback Testing for Interfaces | 461

Set Up BERT Test for Physical Connections | 463

Start and Stop a BERT Test | 464

Locate LINK Alarms and Counters for Gigabit Ethernet | 465

Checklist: Find LINK Alarms and Counters for Ethernet | 466

Display LINK Alarms on Gigabit Ethernet | 466

Gigabit Ethernet Counters | 469

Troubleshoot: 10-Gigabit Ethernet Port Stuck in Down State | 473

Verify Link and Transceivers with Pseudo Random Binary Sequence (PRBS) Test | 475

How to Use Pseudo Random Binary Sequence (PRBS) Test for Link and Transceivers Verification | 475

Configure Unidirectional Diagnostics | 477

Configure Bidirectional Diagnostics with Remote Loopback | 479

Clear the Interface Statistics | 483

9

Configuration Statements and Operational Commands

Junos CLI Reference Overview | 487

About This Guide

Use this guide to configure, monitor, and troubleshoot the various supported Ethernet Interfaces, including aggregated Ethernet Interfaces on Juniper Networks routers.

1

PART

Ethernet Interfaces

- Ethernet Interfaces Overview | **2**
- Configure Ethernet Interfaces | **4**
- Flow Control for Ethernet Interfaces | **6**
- MAC Address Filtering and Accounting on Ethernet Interfaces | **7**
- Management Ethernet Interfaces | **9**
- Point-to-Point Protocol over Ethernet (PPPoE) | **13**

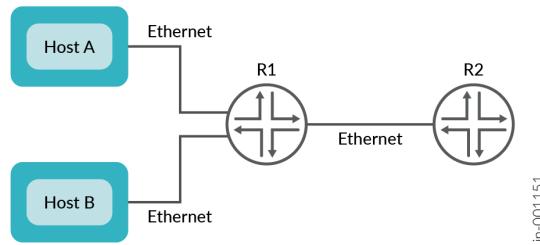
Ethernet Interfaces Overview

SUMMARY

This topic provides an overview of different types of Ethernet interfaces that supports ACX, MX, and PTX routers.

Ethernet Interfaces are networking Interfaces that provide traffic connectivity. You can configure physical Interfaces and logical Interfaces on your device. You can configure the speed of the Interface, limit the rate at which ingress traffic arrives on Fast-Ethernet ports, configure the Interface to operate in full-duplex or half-duplex mode, configure MAC address validation on static Ethernet Interfaces, and other basic configurations.

Figure 1: Ethernet interfaces



Juniper Supported Ethernet Interfaces

Juniper Networks routers support the following types of Ethernet interfaces:

- Tri-Rate Ethernet copper
- Kilobits per second (Kbps)
- Megabits per second (Mbps)
- Gigabit Ethernet
- Gigabit Ethernet intelligent queuing (IQ)

- Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet
- 10-Gigabit Ethernet dense wavelength-division multiplexing (DWDM)
- 25-Gigabit Ethernet
- 40-Gigabit Ethernet
- 100-Gigabit Ethernet
- 200-Gigabit Ethernet
- 400-Gigabit Ethernet
- 800-Gigabit Ethernet
- Management Ethernet interface, which is an out-of-band management interface within the router
- Internal Ethernet interface, which connects the Routing Engine to the packet forwarding components
- Aggregated Ethernet interface, a logical linkage of Gigabit Ethernet, or 10-Gigabit Ethernet physical connections

RELATED DOCUMENTATION

[MAC Address Filtering Overview | 8](#)

Configuring Flow Control

Configuring the Interface Speed on Ethernet Interfaces

Configure Ethernet Interfaces

SUMMARY

Learn how to configure the physical properties of an interface specific to Fast-Ethernet interfaces, Gigabit-Ethernet interfaces, and aggregated Ethernet interfaces.

IN THIS SECTION

- Configure Ethernet Interfaces Properties | 4

Configure Ethernet Interfaces Properties

To configure the Ethernet interfaces properties based on your requirement, select the appropriate command.

1. In configuration mode, go to the interfaces hierarchy level to set the interface port type such as aex, e3, fe, ge, t3.

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

2. Configure the Gigabit Ethernet Interface at the [edit interfaces *ge-fpc/pic/port*] hierarchy level.

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

3. Configure aggregated Ethernet interface, at the [edit interfaces *aex*] hierarchy level

```
[edit interfaces aex]
user@host# set aggregated-ether-options
```

4. Configure the 10-Gigabit Ethernet DWDM interface, include the `optics-options` statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level.

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

5. Configure LAN or WAN on 10-Gigabit Ethernet interface, include the `framing` statement at the `[edit interfaces xe-fpc/pic/port]` hierarchy level.

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

6. Configure OAM 802.3ah support for Ethernet interfaces, include the `oam` statement at the `[edit protocols]` hierarchy level.

```
[edit protocols]
user@host# set oam
```

7. Configure Gigabit Ethernet IQ-specific logical interface, at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level or `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]`

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set input-vlan-map
user@host# set output-vlan-map
user@host# set layer2-policer
user@host# set vlan-tags inner tpid.vlan-id outer tpid.vlan-id
```

8. Configure the speed at the `[edit interfaces interface-name]` hierarchy level.

```
[edit interfaces interface-name]
user@host# set speed (10m | 100m | 1g | auto | auto-10m-100m)
```

9. Configure an Ethernet interface to operate in either full-duplex or half-duplex mode, include the `link-mode` statement at the `[edit interfaces interface-name]` hierarchy level.

```
[edit interfaces interface-name]
user@host# set link-mode (full-duplex | half-duplex)rate
```

10. Configure multicast statistics collection on Ethernet Interfaces, include the `multicast-statistics` statement at the `[edit interfaces interface-name ge-fpc/pic/port]` hierarchy level.

```
[edit interfaces interface-name ge-fpc/pic/port]
user@host# set multicast-statisticsrate
```

11. Configure MAC address validation on static Ethernet interfaces, include the `arp ip-address mac mac-address` statement at the `[edit interfaces interface-name unit logical-unit-number family inet address interface-address]` hierarchy level.

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

Flow Control for Ethernet Interfaces

SUMMARY

Learn about flow control for Ethernet interfaces, how to enable and disable flow control for Ethernet interfaces.

IN THIS SECTION

- [Flow Control Overview | 6](#)

Flow Control Overview

Flow control on an Ethernet interface regulates traffic flow to prevent frames from dropping during network congestion. Flow control stops and resumes the transmission of network traffic between two

connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting prevents node buffers from overflowing and dropping frames.

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic. Ethernet PAUSE also works on aggregated Ethernet interfaces.

Symmetric Flow Control

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface send and respond Ethernet PAUSE messages when flow control is enabled. The interface does not send or respond Ethernet PAUSE messages when flow control is disabled.

To enable symmetric flow control, use the `flow-control` statement at the `[edit interfaces interface-name ether-options]` hierarchy level.

To permit unrestricted traffic, disable the flow control. To disable the flow control, use the `no-flow-control` statement the `[edit interfaces interface-name ether-options]` hierarchy level.

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer.

SEE ALSO

flow-control

MAC Address Filtering and Accounting on Ethernet Interfaces

SUMMARY

Learn how to enable MAC address filtering and how to configure MAC address accounting on Ethernet interfaces.

IN THIS SECTION

 [MAC Address Filtering Overview | 8](#)

- [MAC Address Accounting for Ethernet Interfaces | 9](#)

MAC address filtering is a security feature that controls network access by filtering MAC addresses. To block all incoming packets from a specific MAC address, you can enable MAC address filtering. You can configure an Ethernet Interface to dynamically learn source or destination MAC addresses.

MAC Address Filtering Overview

IN THIS SECTION

- [Configure MAC Address Filtering for Ethernet Interfaces | 8](#)

On Ethernet interfaces with SFPs, you can enable source address filtering to block all incoming packets from a specific MAC address. When you filter logical and physical interfaces, you can specify up to 1000 MAC source addresses per port.

MAC filtering support includes:

- MAC source and destination address filtering for each port.
- MAC source address filtering for each physical interface.
- MAC source address filtering for each logical interface.

Configure MAC Address Filtering for Ethernet Interfaces

On aggregated Ethernet interfaces, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ, and Gigabit Ethernet PICs with SFPs, you can enable source address filtering to block all incoming packets from a specific MAC address.

- To enable the filtering, include the `source-filtering` statement at the `[edit interfaces interface-name aggregated-ether-options | fastether-options | gigether-options]` hierarchy level.
- When source address filtering is enabled, you can configure the interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the `source-address-filter mac-address`

statement at the [edit interfaces *interface-name* aggregated-ether-options | fastether-options | gigether-options] hierarchy levels.

- You can specify the MAC address as *nn:nn:nn:nn:nn:nn* or *nnnn.nnnn.nnnn*, where *n* is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include the source-address-filter statement multiple times.

Source address filtering does not work when Link Aggregation Control Protocol (LACP) is enabled.

If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.

MAC Address Accounting for Ethernet Interfaces

To configure MAC address accounting on an individual Ethernet interface, include the `mac-learn-enable` statement at the [edit interfaces *interface-name* gigether-options ethernet-switch-profile] hierarchy level:

To configure MAC address accounting on an aggregated Ethernet interface, include the `mac-learn-enable` statement at the [edit interfaces *aex* aggregated-ether-options ethernet-switch-profile] hierarchy level.

To prohibit an interface from dynamically learning source and destination MAC addresses, do *not* include the `mac-learn-enable` statement.

To disable dynamic learning of the source and destination MAC addresses after it has been configured, you must delete `mac-learn-enable` from the configuration.

MPCs support MAC address accounting for an individual interface or an aggregated Ethernet interface member link only after the interface has received traffic from the MAC source. If traffic is only exiting an interface, the MAC address is not learned and MAC address accounting does not occur.

Management Ethernet Interfaces

SUMMARY

Learn about the management Ethernet Interface, how to configure the IP address and MAC address on the management Ethernet interfaces.

IN THIS SECTION

 [Management Ethernet Interface Overview | 10](#)

- Configure a Consistent Management IP Address | [11](#)
- Configure the MAC Address on the Management Ethernet Interface | [12](#)
- Platform-Specific Management Ethernet Interfaces Behavior | [12](#)

Management interface in networking is a protocol and format that allows network operators to remotely access, configure, and control network devices. It provides a unified view of the managed device, which helps with consistent management. To connect to the router through the management port, use the management Ethernet interface.

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Review the "[Platform-Specific Management Ethernet Interfaces Behavior](#)" on page [12](#) section for notes related to your platform

Management Ethernet Interface Overview

The router's management Ethernet interface, `fxp0` or `em0`, is an out-of-band management interface that needs to be configured only if you want to connect to the router through the management port on the front of the router. You can configure an IP address and prefix length for this interface, which you commonly do when you first install the Junos OS:

```
[edit]
user@host# set interfaces (fxp0 | em0) unit 0 family inet address/prefix-length
[edit]
user@host# show
interfaces {
    (fxp0 | em0) {
        unit 0 {
            family inet {
                address/prefix-length;
            }
        }
    }
}
```

To determine which management interface type is supported on a router, locate the router and Routing Engine combination in [Supported Routing Engines by Router](#) and note its management Ethernet interface type, either `em0` or `fxp0`.

Alternatively, refer to the Day One + quick start guide for your platform at: [Day One + Guides](#).

Configure a Consistent Management IP Address

On routers with multiple Routing Engines, each Routing Engine is configured with a separate IP address for the management Ethernet interface. To access the primary Routing Engine, you must know which Routing Engine is active and use the appropriate IP address.

Optionally, for consistent access to the primary Routing Engine, you can configure an additional IP address and use this address for the management interface regardless of which Routing Engine is active. This additional IP address is active only on the management Ethernet interface for the primary Routing Engine. During switchover, the address moves to the new primary Routing Engine.

To configure an additional IP address for the management Ethernet interface, include the `master-only` statement at the `[edit groups]` hierarchy level.

In the following example, IP address `10.17.40.131` is configured for both Routing Engines and includes a `master-only` statement. With this configuration, the `10.17.40.131` address is active only on the primary Routing Engine. The address remains consistent regardless of which Routing Engine is active. IP address `10.17.40.132` is assigned to `fxp0` on `re0`, and address `10.17.40.133` is assigned to `fxp0` on `re1`.

```
[edit groups re0 interfaces fxp0]
unit 0 {
    family inet {
        address 10.17.40.131/25 {
            master-only;
        }
        address 10.17.40.132/25;
    }
}
[edit groups re1 interfaces fxp0]
unit 0 {
    family inet {
        address 10.17.40.131/25 {
            master-only;
        }
        address 10.17.40.133/25;
```

```

    }
}
```

Configure the MAC Address on the Management Ethernet Interface

The router's management Ethernet interface uses as its MAC address the MAC address that is burned into the Ethernet card. To change the management Ethernet interface's MAC address, include the `mac` statement at the `[edit interfaces fxp0]` or `[edit interfaces em0]` hierarchy level.

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

To display the MAC address used by the router's management Ethernet interface, enter the `show interface fxp0` or `show interface em0` operational mode command.

SEE ALSO

`mac`

Platform-Specific Management Ethernet Interfaces Behavior

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platform:

Platform	Difference
ACX Series	<ul style="list-style-type: none"> ACX Series routers that support management interface uses <code>em0</code> as the management port for Junos OS platforms. ACX Series routers that support management interface uses <code>re0:mgmt-0</code> as the management port for Junos Evolved platforms.

(Continued)

Platform	Difference
MX Series	<ul style="list-style-type: none"> MX Series routers that support management interface uses <code>fxp0</code> as the management port. The <code>fxp0</code> interface is typically an RJ-45 port on the Routing Engine.
PTX Series	<ul style="list-style-type: none"> PTX Series routers that support management interface support two management interfaces per Routing Engine for Junos Evolved platforms. The two Routing Engines are identified as <code>re0</code> and <code>re1</code>. The RJ-45 copper port on each Routing Engine is indexed as 0, while the SFP fiber port is indexed as 1. <p>For example, the RJ-45 management Ethernet port on Routing Engine 0 is typically named <code>re0:mgmt-0</code>.</p>

Point-to-Point Protocol over Ethernet (PPPoE)

SUMMARY

Learn about PPPoE interface, how to configure, verify, and trace PPPoE on Ethernet interfaces.

IN THIS SECTION

● [PPPoE Overview | 14](#)

Use the Point-to-Point Protocol over Ethernet (PPPoE) encapsulation to connect multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device. This topic provides an overview of PPPoE and explains how to configure PPPoE, verify the configuration, as well as trace PPPoE operations.

PPPoE Overview

IN THIS SECTION

- [PPPoE Interfaces | 14](#)
- [PPPoE Stages | 15](#)
- [Optional CHAP Authentication | 16](#)

The Point-to-Point Protocol over Ethernet (PPPoE) connects multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device. Hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

To use PPPoE, you must configure the router as a PPPoE client, encapsulate PPP packets over Ethernet, and initiate a PPPoE session.

Multiple hosts can be connected to the Services Router, and their data can be authenticated, encrypted, and compressed before the traffic is sent to the PPPoE session on the Services Router's Fast Ethernet or ATM-over-ADSL interface. PPPoE is easy to configure and enables services to be managed on a per-user basis rather than on a per-site basis.

PPPoE Interfaces

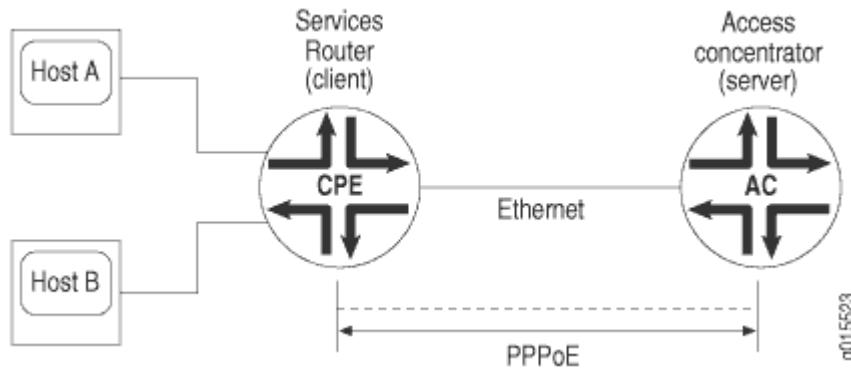
The PPPoE configuration is the same for both interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:

- If the interface is Fast Ethernet, use a PPPoE encapsulation.
- If the interface is ATM over ADSL, use a PPPoE over ATM encapsulation.

Ethernet Interface

The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. [Figure 2 on page 15](#) shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 2: PPPoE Session on an Ethernet Loop



PPPoE Stages

PPPoE has two stages, the discovery stage and the PPPoE session stage. In the discovery stage, the client discovers the access concentrator by identifying the Ethernet media access control (MAC) address of the access concentrator and establishing a PPPoE session ID. In the PPPoE session stage, the client and the access concentrator build a point-to-point connection over Ethernet, based on the information collected in the discovery stage.

If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.

If you do not configure the access concentrator name, the PPPoE session starts using any available server in the network.

PPPoE Discovery Stage

A Services Router initiates the PPPoE discovery stage by broadcasting a PPPoE active discovery initiation (PADI) packet. To provide a point-to-point connection over Ethernet, each PPPoE session must learn the Ethernet MAC address of the access concentrator and establish a session with a unique session ID. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.

A Services Router cannot receive PPPoE packets from two different access concentrators on the same physical interface.

The PPPoE discovery stage consists of the following steps:

1. PPPoE active discovery initiation (PADI)—The client initiates a session by broadcasting a PADI packet on the LAN to request a service.

2. PPPoE active discovery offer (PADO)—Any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client.
3. PPPoE active discovery request (PADR)—From the PADOs it receives, the client selects one access concentrator based on its name or the services offered and sends it a PADR packet to indicate the service or services needed.
4. PPPoE active discovery session-Confirmation (PADS)—When the selected access concentrator receives the PADR packet, it accepts or rejects the PPPoE session.
 - To accept the session, the access concentrator sends the client a PADS packet with a unique session ID for a PPPoE session and a service name that identifies the service under which it accepts the session.
 - To reject the session, the access concentrator sends the client a PADS packet with a service name error and resets the session ID to zero.

PPPoE Session Stage

The PPPoE session stage starts after the PPPoE discovery stage is over. The access concentrator can start the PPPoE session after it sends the PADS packet to the client, or the client can start the PPPoE session after it receives a PADS packet from the access concentrator. A Services Router supports multiple PPPoE sessions on each interface, but no more than 256 PPPoE sessions on all interfaces on the Services Router.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID. After the PPPoE session is established, data is sent as in any other PPP encapsulation. The PPPoE information is encapsulated within an Ethernet frame and is sent to a unicast address. In this stage, both the client and the server must allocate resources for the PPPoE logical interface.

After a session is established, the client or the access concentrator can send a PPPoE active discovery termination (PADT) packet anytime to terminate the session. The PADT packet contains the destination address of the peer and the session ID of the session to be terminated. After this packet is sent, the session is closed to PPPoE traffic.

Optional CHAP Authentication

For interfaces with PPPoE encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP). When you enable CHAP on an interface, the interface can authenticate its peer and be authenticated by its peer.

If you configure an interface to handle incoming CHAP packets only (by including the `passive` statement at the `[edit interfaces interface-name ppp-options chap]` hierarchy level), the interface does not challenge its

peer. However, if the interface is challenged, it responds to the challenge. If you do not include the `passive` statement, the interface always challenges its peer.

2

PART

Aggregated Ethernet Interfaces

- Aggregated Ethernet Interfaces Overview | **19**
- Link Protection of Aggregated Ethernet Interfaces | **42**
- Scheduling on Aggregated Ethernet Interfaces | **46**
- Load Balancing on Aggregated Ethernet Interfaces | **48**
- Performance Monitoring on Aggregated Ethernet Interfaces | **91**
- Periodic Packet Management on Ethernet Interfaces | **93**
- Ethernet Link Aggregation | **95**

Aggregated Ethernet Interfaces Overview

SUMMARY

Learn about aggregated Ethernet interfaces, how to configure an aggregated Ethernet interface, LACP, and other supported features.

IN THIS SECTION

- [Configure Aggregated Ethernet Interfaces | 20](#)
- [Enhanced Link Aggregation Group | 22](#)
- [Mixed-Rate Aggregated Ethernet Interfaces | 23](#)
- [Link Aggregation Control Protocol | 25](#)
- [Configure LACP | 26](#)
- [Targeted Distribution Across Aggregated Ethernet Member Links | 28](#)
- [Configure Targeted Distribution Across Aggregated Ethernet Member Links | 29](#)
- [MAC Address on Aggregated Ethernet Interfaces | 40](#)

You can group or bundle multiple Ethernet interfaces together to form a single link layer interface known as the aggregated Ethernet interface (aex) or a link aggregation group (LAG). The IEEE 802.3ad standard defines link aggregation of Ethernet interfaces and provides a method by which you can group or bundle multiple Ethernet interfaces. Bundling multiple interfaces together enables you to increase the supported bandwidth. The device treats the aggregated Ethernet interface or LAG as a single link instead of a combination of multiple links.

Benefits

- Increased bandwidth and cost effectiveness—The aggregated link provides higher bandwidth than the bandwidth provided by each individual link without requiring new equipment.
- Increased resiliency and availability—if any of the physical links goes down, the traffic is reassigned to another member link.
- Load balancing—The aggregated Ethernet bundle balances the load between its member links if a link fails.

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Review the "Platform-Specific Maximum Interfaces Per LAG Behavior" on page 41 section for notes related to your platform.

Guidelines to Configure Aggregated Ethernet Interfaces

Consider the following guidelines as you configure an aggregated Ethernet interface.

- For Junos OS Evolved, if you add a new member interface to the aggregated Ethernet bundle, a link flap event is generated. The physical interface is deleted as a regular interface and then added back as a member. During this time, the details of the physical interface are lost.
- Use the `gigether-options` statement to configure aggregated Ethernet interfaces on the member link interfaces.

Configure Aggregated Ethernet Interfaces

Follow these steps to configure aggregated Ethernet interfaces on your routing device.

1. Specify the number of aggregated Ethernet bundles you want on your device. If you specify the `device-count` value as 2, you can configure two aggregated bundles.

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

2. Specify the members you want to include within the aggregated Ethernet bundle and add the member individually. Aggregated interfaces are numbered from ae0 through ae4092.

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

3. Specify the link speed for the aggregated Ethernet links. When you specify the speed, all the interfaces that make up the aggregated Ethernet bundle have the same speed. You can also configure the member links of an aggregated Ethernet bundle with a combination of rates—that is, mixed rates—for efficient bandwidth utilization.

```
[edit interfaces]
user@host# set aex aggregated-ether-options link-speed speed
```

4. Specify the minimum number of links for the aggregated Ethernet interface (aex) —that is, the defined bundle— to be labeled *up*. By default, only one link must be up for the bundle to be labeled *up*.

```
[edit interfaces]
user@host# set aex aggregated-ether-options minimum-links number
```

You cannot configure the minimum number of links and the minimum bandwidth at the same time. They are mutually exclusive.

5. (Optional) Specify the minimum bandwidth for the aggregated Ethernet links. You cannot configure link protection with minimum bandwidth.

```
[edit interfaces]
user@host# set aex aggregated-ether-options minimum-bandwidth
```

6. Specify an interface family and the IP address for the aggregated Ethernet bundle. Aggregated Ethernet interfaces can be VLAN-tagged or untagged.

```
[edit interfaces]
user@host# set aex vlan-tagging unit 0 vlan-id vlan-id
```

Untagged Interface

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

7. (Optional) Configure your device to collect multicast statistics for the aggregated Ethernet interface.

```
[edit interfaces]
user@host# set aex multicast-statistics
```

8. Verify and commit the configuration.

```
[edit interfaces]
user@host# run show configuration
user@host# commit
```

9. (Optional) Delete an aggregated Ethernet Interface.

```
[edit]
user@host# delete interfaces aex
```

SEE ALSO

minimum-bandwidth

minimum-links

Enhanced Link Aggregation Group

IN THIS SECTION

- [Benefits | 22](#)

When you associate a physical interface with an aggregated Ethernet interface, the physical child links are also associated with the parent aggregated Ethernet interface to form a link aggregation group (LAG). So, one child next hop is created for each member link of an aggregated Ethernet interface for each VLAN interface. For example, an aggregate next hop for an aggregated Ethernet interface with 16 member links leads to the creation of 17 next hops per VLAN.

When you configure enhanced LAG, child next hops are not created for member links and, as a result, a higher number of next hops can be supported. To configure enhanced LAG, you must configure the device's network services mode as `enhanced-ip`. This feature is not supported if the device's network services mode is set to operate in the `enhanced-ethernet` mode. This feature is enabled by default if the network services mode on the device is configured as `enhanced-mode`.

Benefits

- Reduction in memory and CPU usage to support aggregated Ethernet interfaces.
- Improvement in system performance and scaling numbers.

Mixed-Rate Aggregated Ethernet Interfaces

IN THIS SECTION

- [Benefits | 23](#)

On Juniper Networks devices, you can configure the member links of an aggregated Ethernet bundle to operate at different link speeds (also known as rates). The configured aggregated Ethernet bundle is known as a mixed-rate *aggregated Ethernet bundle*. When you configure the member links of an aggregated Ethernet bundle in LAN mode as well as WAN mode for 10-Gigabit Ethernet interfaces, the configuration is known as mixed-mode configuration.

Benefits

- Efficient bandwidth utilization—When you configure the member links with different link speeds, the bandwidth is efficiently and completed used.
- Load balancing—Balances the load between member links within an aggregated Ethernet bundle if a link fails.

Guidelines to Configure Mixed-Rate Aggregated Ethernet Links

Consider the following guidelines as you configure a mixed-rate aggregated Ethernet bundle:

- You can configure a maximum of 64 member links to form a mixed aggregated Ethernet bundle.
- When you mix a 10-Gigabit Ethernet interface in LAN mode and a 10-Gigabit Ethernet interface in WAN mode in the same aggregated bundle on the routers, it is not considered a mixed-rate aggregate. To mix the interfaces having the same speed but different framing options, you need not use the `mixed` statement at the `[edit interfaces interface-name aggregated-ether-options link-speed]` hierarchy level.
- Mixed-rate aggregated Ethernet links can interoperate with non-Juniper Networks aggregated Ethernet member links provided that mixed-rate aggregated Ethernet load balancing is configured at egress.
- After you configure a mixed-rate aggregated Ethernet link on a 100-Gigabit Ethernet PIC with CFP, changing aggregated Ethernet link protection or LACP link protection configurations results in aggregated Ethernet link flapping. Also, changing the configuration of a mixed aggregated Ethernet link can result in aggregated Ethernet link flapping.

- If a member link fails, fast reroute is activated. For mixed-mode AE, the repair occurs by rebalancing the weighting array. This rebalancing considers the weights of the currently active links into account. The weights remain proportional to the link speed of the member.
- Mixed-rate aggregated Ethernet links do not support rate-based CoS components such as scheduler, shaper, and policer. However, the default CoS settings are supported on the mixed-rate aggregated Ethernet links.
- Load balancing of the egress traffic across the member links of a mixed-rate aggregated Ethernet link is proportional to the rates of the member links. Egress multicast load balancing is not supported on mixed aggregated Ethernet interfaces.
- Mixed-rate aggregated Ethernet interface do not support aggregated Ethernet link protection, link protection on a 1:1 model, and LACP link protection.

Configure Mixed-Rate Aggregated Ethernet Interfaces

Follow these steps to configure mixed-rate aggregated Ethernet interfaces on your routing device.

1. Specify the number of aggregated Ethernet bundles you want on your device. If you specify the device-count value as 2, you can configure two aggregated bundles.

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

2. Specify the members you want to include within the aggregated Ethernet bundle and add the member individually. Aggregated interfaces are numbered from ae0 through ae4092.

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

3. Specify the link speed for the aggregated Ethernet links. When you specify the speed as mixed, you can configure the member links of an aggregated Ethernet bundle with a combination of rates—that is, mixed rates—for efficient bandwidth utilization.

```
[edit interfaces]
user@host# set aex aggregated-ether-options link-speed mixed
```

4. Specify the minimum bandwidth for the aggregated Ethernet links.

```
[edit interfaces]
user@host# set aex aggregated-ether-options minimum-bandwidth
```

5. Verify and commit the configuration.

```
[edit interfaces]
user@host# run show configuration
user@host# commit
```

SEE ALSO

| *minimum-bandwidth*

Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP), defined in IEEE 802.3ad, is a monitoring protocol that detects link-layer failure within a network. You can use LACP to monitor the local and remote ends of member links in a LAG.

By default, LACP is not configured on aggregated Ethernet interfaces. Ethernet links do not exchange information about the state of the link. When you configure LACP, the transmitting link (also known as *actor*) initiates transmission of LACP packets to the receiving link (also known as *partner*). The actor is the local interface in a LACP exchange. The partner is the remote interface in a LACP exchange.

When you configure LACP, you must select one of the following transmission modes for each end of the LAG:

- Active-To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode. If either the actor or partner is active, they exchange LACP packets.
- Passive-No exchange of LACP packets. This is the default transmission mode.

Benefits

- Link monitoring—LACP detects invalid configurations on the local end as well as the remote end of the link.

- Link resiliency and redundancy—If a link fails, LACP ensures that traffic continues to flow on the remaining links.

Guidelines to Configure LACP

Consider the following guidelines when you configure LACP:

- When you configure LACP on multiple different physical interfaces, only features that are supported across all of the linked devices are supported in the resulting link aggregation group (LAG) bundle. For example, different PICs can support a different number of forwarding classes. When you link ports of a 16-forwarding-class PIC with an 8-forwarding-class PIC using link aggregation, the resulting LAG bundle supports up to 8 forwarding classes. Similarly, linking together a PIC that supports weighted random early detection (WRED) with a PIC that does not support it results in a LAG bundle that does not support WRED.
- If you configure the LACP system identifier (by using the `system-id systemid` statement) to be all zeros (00:00:00:00:00:00), the commit operation throws an error.
- If the aggregated Ethernet bundle is up, you can enable a device to handle packets on a member link, irrespective of the LACP state. This is done using the `accept-data` statement. But, this method doesn't comply with the packet processing that is defined in the IEEE 802.3ax standard. According to this standard, the packets should be dropped, but they are processed instead because you configured the `accept-data` statement.

Configure LACP

Follow these steps to configure LACP on an aggregated Ethernet interfaces on your routing device.

1. Specify the LACP transmission mode - active or passive

```
[edit interfaces interface-name aggregated-ether-options]
user@host# set lacp active
user@host# set lacp passive
```

- Specify the interval at which the interfaces send LACP packets. When you configure different intervals for the active and passive interfaces, the *actor* transmits the packets at the rate that is configured on the *partner's* interface.

```
[edit interfaces interface-name aggregated-ether-options lacp]
user@host# set periodic interval
```

- Configure the LACP system ID. The user-defined system identifier in LACP enables two ports from two different devices to act as though they were part of the same aggregate group. The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

```
[edit interfaces interface-name aggregated-ether-options lacp]
user@host# set system-id system-id
```

- Configure the LACP system priority at the Aggregated Ethernet interface level. This system priority takes precedence over the priority value that is configured at the global [edit chassis] level. The device with numerically lower value (higher priority value) becomes the controlling device. If both devices have the same LACP system priority value, the device MAC address determines which device is in control.

```
[edit interfaces interface-name aggregated-ether-options lacp]
user@host# set system-priority system-priority
```

- (Optional) Configure the LACP administrative key.

```
[edit interfaces interface-name aggregated-ether-options lacp]
user@host# set admin-key number
```

- Specify the time period, in seconds, for which LACP maintains the state of a member link as expired. To prevent excessive flapping of a LAG member link, you can configure LACP to prevent the transition of an interface from down to up for a specified interval.

```
[edit interfaces interface-name aggregated-ether-options lacp]
user@host# set hold-time timer-value
```

7. Configure the device to process packets received on a member link irrespective of the LACP state if the aggregated interface status is up.

```
[edit interfaces interface-name aggregated-ether-options lacp]
user@host# set accept-data
```

8. Verify and commit the configuration.

```
[edit interfaces interface-name aggregated-ether-options lacp]
user@host# run show configuration
user@host# commit
```

SEE ALSO

hold-time

periodic

system-id

system-priority

Targeted Distribution Across Aggregated Ethernet Member Links

Aggregated Ethernet bundles use a hash-based algorithm to distribute traffic over multiple links. Traffic destined through a logical interface of a bundle can exit through any of the member links that are based on the hashing algorithm. Egress policy is distributed between individual member interface schedulers or policers that are instantiated in each Packet Forwarding Engine hosting a member link. Distributed egress policy enforcement relies on traffic load balancing and so is not always accurate.

Targeted distribution provides a mechanism to direct traffic through specified links of an aggregated Ethernet bundle. You can also use targeted distribution to assign roles to member links to handle link failure scenarios. Targeted distribution ensures accurate policy enforcement that is not distributed for a given logical interface. Targeted distribution is applicable to both Layer 2 and Layer 3 interfaces, irrespective of the family configured for the logical interface. The outbound traffic of a Layer 3 host is distributed among all the member links of an aggregated Ethernet bundle. Targeted distribution is implemented only for the transit traffic.

You can form distribution lists consisting of member links of the aggregated Ethernet interfaces and you can assign roles to these lists, as follows:

- Primary distribution list: You can configure the member links that will be part of the primary distribution list. Traffic is load-balanced among all the member links in the primary list. If all links within the primary list are up, traffic is forwarded on those links. If some of the links within a primary list fail, the remaining links carry traffic.
- Backup distribution list: You can configure the member links that will be part of the backup distribution list. If all links within the primary list go down, only then the links in the backup list start carrying traffic. If some of links within the backup list fail, the remaining links in the backup list carry traffic.
- Standby distribution list: All remaining links are added to the defined standby list. If all the links within the primary list and the backup list go down, only then the links in the standby list start carrying traffic. When the links in the primary distribution list come back online, they resume carrying traffic.

Benefits

- Accurate policy enforcement—Policy enforcement is not distributed and is, therefore, accurate.
- Load balancing—With targeted distribution, you can load-balance the traffic between the aggregated Ethernet bundle member links.

Configure Targeted Distribution Across Aggregated Ethernet Member Links

IN THIS SECTION

- Requirements | [33](#)
- Overview | [33](#)
- Verification | [34](#)

This example shows how to configure primary and backup targeted distribution lists for aggregated Ethernet member links. Member links are assigned membership to the distribution lists. Logical interfaces of the aggregated Ethernet bundle are then assigned membership to the primary list and the backup list.

Configuration

CLI Quick Configuration

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

```
[edit groups GR-AE-ACCESS-DISTRIBUTION]
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description "matched-odd" targeted-
distribution primary-list dl2
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description "matched-odd" targeted-
distribution backup-list dl1
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description "matched-even" targeted-
distribution primary-list dl1
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description "matched-even" targeted-
distribution backup-list dl2
user@host# set interfaces <ae*> apply-groups GR-AE-ACCESS-DISTRIBUTION
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation flexible-ethernet-services
unit 101 vlan-id 101 family inet address 10.1.0.1/16
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation flexible-ethernet-services
unit 102 vlan-id 102 family inet address 10.2.0.1/16
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation flexible-ethernet-services
unit 103 vlan-id 103 family inet address 10.3.0.1/16
user@host# set interfaces <ae*> flexible-vlan-tagging encapsulation flexible-ethernet-services
unit 104 vlan-id 104 family inet address 10.4.0.1/16
```

Step-by-Step Procedure

To configure targeted distribution:

1. Create a global apply group and specify the primary list and the backup list.

```
[edit groups GR-AE-ACCESS-DISTRIBUTION]
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description "matched-odd" targeted-
distribution primary-list dl2
user@host# set interfaces <ae*> unit <*[1 3 5 7 9]> description "matched-odd" targeted-
distribution backup-list dl1
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description "matched-even" targeted-
distribution primary-list dl1
```

```
user@host# set interfaces <ae*> unit <*[0 2 4 6 8]> description "matched-even" targeted-distribution backup-list dl2
```

2. Attach the defined apply group to the aggregated Ethernet interface.

```
[edit]
user@host# set interfaces ae10 apply-groups GR-AE-ACCESS-DISTRIBUTION
```

3. Create the logical interfaces and configure its parameters.

```
[edit]
user@host# set interfaces ae10 apply-groups GR-AE-ACCESS-DISTRIBUTION
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation flexible-ethernet-services
  set unit 101 vlan-id 101 family inet address 10.1.0.1/16
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation flexible-ethernet-services
  unit 102 vlan-id 102 family inet address 10.2.0.1/16
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation flexible-ethernet-services
  unit 103 vlan-id 103 family inet address 10.3.0.1/16
user@host# set interfaces ae10 flexible-vlan-tagging encapsulation flexible-ethernet-services
  unit 104 vlan-id 104 family inet address 10.4.0.1/16
```

Results

From configuration mode, confirm your configuration by using the `show` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show groups GR-AE-ACCESS-DISTRIBUTION
interfaces {
  <ae*> {
    unit "<*[1 3 5 7 9]>" {
      description "matched odd";
      targeted-distribution {
        primary-list dl2;
        backup-list dl1;
      }
    }
    unit "<*[0 2 4 6 8]>" {
      description "matched even";
      targeted-distribution {
```

```
        primary-list dl1;
        backup-list dl2;
    }
}
}
}
```

```
user@host# show interfaces ae10 apply-groups
apply-groups GR-AE-ACCESS-DISTRIBUTION;
```

```
user@host# show interfaces ae10
apply-groups GR-AE-ACCESS-DISTRIBUTION;
flexible-vlan-tagging; encapsulation flexible-ethernet-services;
unit 101 {
    vlan-id 101;
    family inet {
        address 10.1.0.1/16 {
        }
    }
}
unit 102 {
    vlan-id 102;
    family inet {
        address 10.2.0.1/16 {
        }
    }
}
unit 103 {
    vlan-id 103;
    family inet {
        address 10.3.0.1/16 {
        }
    }
}
unit 104 {
    vlan-id 104;
    family inet {
        address 10.4.0.1/16 {
        }
    }
}
```

```
    }  
}
```

Requirements

This example uses the following software and hardware components:

- Any suggested Junos releases
- One MX Series 5G Universal Routing Platform

Overview

Targeted distribution provides a mechanism to direct traffic through specified links of an aggregated Ethernet bundle, and also assigns roles to member links to handle link failure scenarios. You can configure targeted distribution to load-balance the traffic between the aggregated Ethernet bundle member links. You can map a logical interface to a single link only for the outgoing traffic.

This example uses the `apply-groups` configuration for specifying the distribution lists for the logical interfaces of the aggregated Ethernet member links. You can use the `apply-groups` statement to inherit the Junos OS configuration statements from a configuration group. The `apply-groups` configuration statement in the example shows the odd-numbered member links of the aggregated Ethernet bundle being assigned the primary list `dl2` and even-numbered member links being assigned primary list `dl1`.

The aggregated Ethernet interface used in this example is `ae10` with units 101, 102, 103, and 104. The physical interface `ge-0/0/3` is specified as distribution list `dl1` and `ge-0/0/4` as `dl2`. The distribution list `dl1` is assigned as the primary list for those logical interface unit numbers of the aggregated Ethernet bundle ending in an odd number. Alternatively, the distribution list `dl2` is the primary list for those ending in an even number.

To configure targeted distribution, you must:

1. Create a global apply group.
2. Assign each member of the aggregated Ethernet interface to a different distribution list.
3. Attach the apply group to the aggregated Ethernet interface.
4. Create the logical interfaces. The apply group automatically assigns the distribution lists to each member of the aggregated Ethernet bundle as required.

Verification

IN THIS SECTION

- [Verify Targeted Distribution of Logical Interfaces | 34](#)

Verify Targeted Distribution of Logical Interfaces

Purpose

Verify that the logical interfaces are assigned to the distribution lists.

Action

To verify that the logical interfaces are assigned to the distribution lists, enter the `show interfaces detail` or `extensive` command.

The `show interfaces detail` or `extensive` command output shows the logical interfaces ending in an odd number being assigned to the distribution list `d1`(`ge-0/0/3`) and those ending in an even number being assigned to the distribution list `d2`(`ge-0/0/4`) by default. If there is a failure of either of those interfaces, the logical interfaces switch to the interfaces in the backup list or continue to use the active member interface. For example, on the aggregated Ethernet bundle `ae10.101`, the primary interface shown is `ge-0/0/4` and on the aggregated Ethernet bundle `ae10.102`, the primary interface is `ge-0/0/3`, and similarly for the other logical interfaces.

```
user@host# run show interfaces extensive ae10

Physical interface: ae10, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 612, Generation: 132
  Link-level type: Flexible-Ethernet, MTU: 9000, Speed: 2Gbps, BPDU Error: None, MAC-REWRITE
  Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 1, Minimum bandwidth needed: 1bps
  Device flags  : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:05:86:1e:70:c1, Hardware address: 00:05:86:1e:70:c1
  Last flapped   : 2016-08-30 16:15:28 PDT (00:43:15 ago)
```

```

Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 77194 200 bps
  Input packets: 0 0 pps
  Output packets: 300 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Dropped traffic statistics due to STP State:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
  errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
  0 0 0 0
  1 0 0 0
  2 0 0 0
  3 0 0 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
  0 0 0 0
  1 0 0 0
  2 0 0 0
  3 0 0 0
Queue number: Mapped forwarding classes
  0 best-effort
  1 expedited-forwarding
  2 assured-forwarding
  3 network-control

Logical interface ae10.101 (Index 345) (SNMP ifIndex 617) (Generation 154)
  Description: matched odd
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.101 ] Encapsulation: ENET2
  Statistics  Packets  pps  Bytes  bps

```

```

Bundle:
  Input : 0 0 0 0
  Output: 2 0 92 0

Adaptive Statistics:
  Adaptive Adjusts: 0
  Adaptive Scans : 0
  Adaptive Updates: 0

Link:
  ge-0/0/3.101
    Input : 0 0 0 0
    Output: 2 0 92 0
  ge-0/0/4.101
    Input : 0 0 0 0
    Output: 0 0 0 0

```

Aggregate member links: 2

Marker Statistics:	Marker Rx	Resp Tx	Unknown Rx	Illegal Rx
ge-0/0/3.101	0	0	0	0
ge-0/0/4.101	0	0	0	0

List-Type	Status
Primary	Active
Interfaces:	
ge-0/0/4	Up
List-Type	Status
Backup	Waiting
Interfaces:	
ge-0/0/3	Up
List-Type	Status
Standby	Down

```

Protocol inet, MTU: 8978, Generation: 198, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.1.0.1/15, Local: 10.1.0.2, Broadcast: 10.1.0.3, Generation: 154
Protocol multiservice, MTU: Unlimited, Generation: 199, Route table: 0
  Policer: Input: __default_arp_policer__

```

Logical interface ae10.102 (Index 344) (SNMP ifIndex 615) (Generation 153)

Description: matched even

Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.102] Encapsulation: ENET2

Statistics	packets	pps	bytes	bps	
Bundle:					
Input :	0	0	0	0	
Output:	4	0	296	0	
Adaptive Statistics:					
Adaptive Adjusts:	0				
Adaptive Scans :	0				
Adaptive Updates:	0				
Link:					
ge-0/0/3.102					
Input :	0	0	0	0	
Output:	4	0	296	0	
ge-0/0/4.102					
Input :	0	0	0	0	
Output:	0	0	0	0	
Marker Statistics:		Marker Rx	Resp Tx	Unknown Rx	Illegal Rx
ge-0/0/3.102		0	0	0	0
ge-0/0/4.102		0	0	0	0
List-Type Status					
Primary	Active				
Interfaces:					
ge-0/0/3		Up			
List-Type Status					
Backup	Waiting				
Interfaces:					
ge-0/0/4		Up			
List-Type Status					
Standby	Down				
Protocol inet, MTU: 8978, Generation: 196, Route table: 0					
Flags: Sendbcast-pkt-to-re					
Addresses, Flags: Is-Preferred Is-Primary					
Destination: 10.2.0.1 , Local: 10.2.0.1, Broadcast: 10.2.0.3, Generation: 152					
Protocol multiservice, MTU: Unlimited, Generation: 197, Route table: 0					
Policer: Input: __default_arp_policer__					
Logical interface ae10.103 (Index 343) (SNMP ifIndex 614) (Generation 152)					
Description: matched odd					
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.103] Encapsulation: ENET2					
Statistics	packets	pps	bytes	bps	
Bundle:					
Input :	0	0	0	0	

```

Output: 3 0 194 0
Adaptive Statistics:
  Adaptive Adjusts: 0
  Adaptive Scans : 0
  Adaptive Updates: 0
Link:
  ge-0/0/3.103
    Input : 0 0 0 0
    Output: 3 0 194 0
  ge-0/0/4.103
    Input : 0 0 0 0
    Output: 0 0 0 0
Marker Statistics:  Marker Rx  Resp Tx  Unknown Rx  Illegal Rx
  ge-0/0/3.103  0 0 0 0
  ge-0/0/4.103  0 0 0 0

List-Type Status
Primary Active
  Interfaces:
    ge-0/0/4 Up
List-Type Status
Backup Waiting
  Interfaces:
    ge-0/0/3 Up
List-Type Status
Standby Down

Protocol inet, MTU: 8978, Generation: 194, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.3.0.0/15, Local: 10.3.0.1, Broadcast: 10.3.0.3, Generation: 150
Protocol multiservice, MTU: Unlimited, Generation: 195, Route table: 0
  Policer: Input: __default_arp_policer__

Logical interface ae10.104 (Index 342) (SNMP ifIndex 616) (Generation 151)
  Description: matched even
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.104 ] Encapsulation: ENET2
  Statistics  Packets      pps      Bytes      bps
  Bundle:
    Input : 0 0 0 0
    Output: 2 0 92 0
  Adaptive Statistics:
    Adaptive Adjusts: 0

```

```

        Adaptive Scans  :      0
        Adaptive Updates:      0
Link:
  ge-0/0/3.104
    Input :      0      0      0      0
    Output:      2      0     92      0
  ge-0/0/4.104
    Input :      0      0      0      0
    Output:      0      0      0      0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
  ge-0/0/3.104          0          0          0          0
  ge-0/0/4.104          0          0          0          0

List-Type      Status
Primary       Active
  Interfaces:
    ge-0/0/3      Up
List-Type      Status
Backup        Waiting
  Interfaces:
    ge-0/0/4      Up
List-Type      Status
Standby       Down

Protocol inet, MTU: 8978, Generation: 192, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.4.0.0/16, Local: 10.4.0.1, Broadcast: 10.4.0.3, Generation: 148
Protocol multiservice, MTU: Unlimited, Generation: 193, Route table: 0
  Policer: Input: __default_arp_policer__

Logical interface ae10.32767 (Index 341) (SNMP ifIndex 613) (Generation 150)
  Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
  Statistics      Packets      pps      Bytes      bps
  Bundle:
    Input :      0      0      0      0
    Output:      0      0      0      0
  Adaptive Statistics:
    Adaptive Adjusts:      0
    Adaptive Scans  :      0
    Adaptive Updates:      0
Link:

```

```

ge-0/0/3.32767
  Input :      0      0      0      0
  Output:    95      0  38039      0
ge-0/0/4.32767
  Input :      0      0      0      0
  Output:    95      0  38039      0
Marker Statistics:  Marker Rx  Resp Tx  Unknown Rx  Illegal Rx
  ge-0/0/3.32767      0      0      0      0
  ge-0/0/4.32767      0      0      0      0
Protocol multiservice, MTU: Unlimited, Generation: 191, Route table: 0
Flags: None
Policer: Input: __default_arp_policer__

```

MAC Address on Aggregated Ethernet Interfaces

You can configure source MAC address and destination MAC address-based accounting for MAC addresses that are dynamically learned on aggregated Ethernet interfaces.

By default, dynamic learning of source and destination MAC addresses on aggregated Ethernet interfaces is disabled. When you enable this feature, you can configure source and destination MAC address-based accounting on the routed interfaces on MX Series routers with DPCs and MPCs. Also, when you enable dynamic learning of MAC addresses, the MAC-filter settings for each member link of the aggregated Ethernet bundle is updated. The limit on the maximum number of MAC addresses that can be learned from an interface does not apply to this dynamic learning of MAC addresses functionality.

Destination MAC-based accounting is supported only for MAC addresses dynamically learned at the ingress interface, including each individual child or member link of the aggregated Ethernet bundle. MPCs do not support destination MAC address learning. Dynamic learning of MAC addresses can be supported on only the aggregated Ethernet interface or on selective individual member links. MAC learning support on the bundle depends on the capability of individual member links. If a link in the bundle does not contain the capability to support MAC learning or accounting, it is disabled on the aggregated Ethernet bundle.

The MAC data for the aggregated bundle is displayed after collecting data from individual child links. On DPCs, these packets are accounted in the egress direction (Output Packet/Byte count), whereas on MPCs, these packets are not accounted because DMAC learning is not supported. This difference in behavior also occurs between child links on DPCs and MPCs. Because this feature to enable dynamic learning is related to collecting MAC database statistics from child links that are based on the command

that is issued from the CLI, there is an impact on the time it takes to display the data on the console based on the size of the MAC database and the number of child links spread across different FPCs.

Benefits

Compute Statistics—Enables you to compute MAC Address statistics for dynamically learned MAC addresses.

Platform-Specific Maximum Interfaces Per LAG Behavior

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platform:

Platform	Difference
ACX Series	<p>ACX7000 series routers from Junos OS Evolved 24.1R1 release support Aggregated Ethernet (AE) Hierarchical Quality of Service (HQoS). The system includes two distinct interface pools: pool_0 and pool_1.</p> <ul style="list-style-type: none"> • ae0 to ae2000 when configured will use pool_0 each interface accommodates up to 16 members • ae2001 to ae4091 when configured will use pool_1 each interface accommodates up to 256 members
MX Series	MX Series routers that support LAG, support a maximum of 64 interfaces per LAG.
PTX Series	PTX Series routers that support LAG, support a maximum of 64 interfaces per LAG.

Link Protection of Aggregated Ethernet Interfaces

SUMMARY

Learn how to provide link protection for aggregated Ethernet Interfaces and configure the minimum number of links in an aggregated Ethernet interfaces bundle.

IN THIS SECTION

- [Aggregated Ethernet Link Protection | 42](#)
- [Configure Aggregated Ethernet Minimum Links | 44](#)
- [Example: Configure Aggregated Ethernet Link Protection | 45](#)

Aggregated Ethernet Link Protection

IN THIS SECTION

- [Configure Link Protection for Aggregated Ethernet Interfaces | 43](#)
- [Configure Primary and Backup Links for Link Aggregated Ethernet Interfaces | 43](#)
- [Revert Traffic on Aggregate Ethernet Interfaces to a Primary Link When Traffic is Passing Through a Backup Link | 43](#)
- [Disable Link Protection for Aggregated Ethernet Interfaces | 44](#)

You can configure link protection for aggregated Ethernet interfaces to provide QoS on the links during operation.

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

ACX7000 Series Routers support revertive mode of link protection. You can enable auto-revertive operation using the following command:

```
set interfaces ae1 aggregated-ether-options link-protection revertive
```

Configure Link Protection for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces support link protection to ensure QoS on the interface.

To configure link protection:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

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

2. Configure the link protection mode.

```
[edit interfaces aex aggregated-ether-options]
user@host# set link-protection
```

Configure Primary and Backup Links for Link Aggregated Ethernet Interfaces

To configure link protection, you must specify a primary and a secondary, or backup, link.

To configure a primary link and a backup link:

1. Configure the primary logical interface.

```
[edit interfaces interface-name]
user@host# set (fastether-options | gigether-options) 802.3ad aex primary
```

2. Configure the backup logical interface.

```
[edit interfaces interface-name]
user@host# set (fastether-options | gigether-options) 802.3ad aex backup
```

Revert Traffic on Aggregate Ethernet Interfaces to a Primary Link When Traffic is Passing Through a Backup Link

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

To manually control when traffic should be diverted back to the primary link from the designated backup link, enter the following operational command:

```
user@host> request interface revert aex
```

SEE ALSO

request interface (revert / switchover) (Aggregated Ethernet Link Protection)

Disable Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the `delete interfaces aex aggregated-ether-options link-protection` configuration command.

```
user@host# delete interfaces aex aggregated-ether-options link-protection
```

SEE ALSO

request interface (revert / switchover) (Aggregated Ethernet Link Protection)

Configure Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled up. By default, only one link must be up for the bundle to be labeled up.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

```
user@host# edit interfaces interface-name aggregated-ether-options
```

2. Configure the minimum number of links.

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

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

- The specified minimum number of links are up.
- The links are in *collecting distributing* state—that is, collecting and distributing states are merged together to form a combined state (coupled control) for the aggregated port. Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the aggregated-ether-options statement, the configuration commit fails and an error message is displayed.

Example: Configure Aggregated Ethernet Link Protection

The following configuration enables link protection on the ae0 interface, and specifies the ge-1/0/0 interface as the primary link and ge-1/0/1 as the secondary link.

```
[edit interfaces]
ae0 {
    aggregated-ether-options {
        link-protection;
    }
}
[edit interfaces]
ge-1/0/0 {
    gigether-options {
        802.3ad ae0 primary;
    }
}
[edit interfaces]
ge-1/0/1 {
    gigether-options {
```

```
802.3ad ae0 backup;
}
}
```

Scheduling on Aggregated Ethernet Interfaces

SUMMARY

Learn how to configure shared scheduling on aggregated Ethernet interfaces in link-protection mode or without link protection.

IN THIS SECTION

- [Configure Shared Scheduling on Aggregated Ethernet Interfaces | 46](#)
- [Configure Scheduler on Aggregated Ethernet Interfaces Without Link Protection | 47](#)

Scheduling on an aggregate interface allows you to configure class-of-service parameters, such as shaping and queuing, for the interface. You can configure these parameters for per-unit, hierarchical, or physical and logical interface schedulers.

Configure Shared Scheduling on Aggregated Ethernet Interfaces

You can configure shared scheduling on aggregated Ethernet interfaces in link protection mode on Gigabit Ethernet Intelligent Queuing 2 (IQ2) and Ethernet Enhanced IQ2 (IQ2E) PICs on routers.

To configure shared scheduling on aggregated Ethernet interfaces:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

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

2. Configure the link protection mode.

```
[edit interfaces aex aggregated-ether-options]
user@host# set link-protection
```

3. Configure shared scheduling.

```
[edit interfaces ae1 aggregated-ether-options]
user@host# top
[edit]
user@host# edit interfaces ae1 shared-scheduler
```

Configure Scheduler on Aggregated Ethernet Interfaces Without Link Protection

The scheduler functions supported are:

- Per unit scheduler
- Hierarchical scheduler
- Shaping at the physical interface

In link-protect mode, only one link is active at a time and the other link acts as the backup link, whereas in a non link-protect mode, all the links of the aggregate bundle are active at the same time. There is no backup link. If a link goes down or a new link is added to the bundle, traffic redistribution occurs.

To configure the hierarchical scheduler on aggregated Ethernet interfaces in the non link-protect mode, include the `hierarchical-scheduler` statement at the `[edit interfaces ae1]` hierarchy level:

```
[edit interfaces ae1 hierarchical-scheduler]
```

To specify the member link bandwidth derivation based on the equal division model (`scale`) or the replication model (`replicate`) on aggregated Ethernet interfaces, include the `member-link-scheduler (scale | replicate)` option at the `[edit class-of-service interfaces ae1]` hierarchy level. The default setting is `scale`.

```
[edit class-of-service interfaces ae1 member-link-scheduler (scale | replicate)]
```

Load Balancing on Aggregated Ethernet Interfaces

SUMMARY

Load balancing on aggregated ethernet interfaces reduces network congestion by dividing traffic among multiple interfaces.

IN THIS SECTION

- Load Balancing and Ethernet Link Aggregation Overview | [49](#)
- Understanding Aggregated Ethernet Load Balancing | [49](#)
- Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data | [51](#)
- Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | [54](#)
- Configuring Adaptive Load Balancing | [56](#)
- Understanding Symmetric Hashing for Load Balancing | [57](#)
- Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers | [58](#)
- Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers | [66](#)
- Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers | [68](#)
- Example: Configuring Aggregated Ethernet Load Balancing | [71](#)
- Platform-Specific Aggregated Ethernet Load Balancing Behavior | [89](#)

When you bundle several physical aggregated Ethernet Interfaces to form a single logical interface, it is called link aggregation. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, increases availability and provides load-balancing capabilities. Load balancing enables the device to divide incoming and outgoing traffic along multiple interfaces to reduce congestion in the network. This topic describes load balancing and how to configure load balancing on your device.

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Review the ["Platform-Specific Aggregated Ethernet Load Balancing Behavior" on page 89](#) section for notes related to your platform.

Load Balancing and Ethernet Link Aggregation Overview

You can create a LAG for a group of Ethernet ports. L2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. Each LAG bundle contains up to 16 links. Platform support depends on the Junos OS release in your installation.

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. The hash-mode of the hashing algorithm is set to L2 payload by default. When the hash-mode is set to L2 payload, the hashing algorithm uses the IPv4 and IPv6 payload fields for hashing. You can also configure the load balancing hash key for L2 traffic to use fields in the L3 and Layer 4 headers using the payload statement. However, note that the load-balancing behavior is platform-specific and based on appropriate hash-key configurations.

For more information, see *Configuring Load Balancing on a LAG Link*. In an L2 switch, one link is overutilized and other links are underutilized.

Understanding Aggregated Ethernet Load Balancing

The link aggregation feature is used to bundle several physical aggregated Ethernet interfaces to form one logical interface. One or more links are aggregated to form a virtual link or link aggregation group (LAG). The MAC client treats this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

In addition to these benefits, an aggregated Ethernet bundle is enhanced to provide load-balancing capabilities that ensure that the link utilization among the member links of the aggregated Ethernet bundle are fully and efficiently utilized.

The load-balancing feature allows a device to divide incoming and outgoing traffic along multiple paths or interfaces in order to reduce congestion in the network. Load balancing improves the utilization of various network paths and provides more effective network bandwidth.

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

Aggregated Interfaces (also called AE for aggregated Ethernet, and AS for aggregated SONET) are a Layer 2 mechanism for load-balancing across multiple interfaces between two devices. Because this is a Layer 2 load-balancing mechanism, all of the individual component links must be between the same two devices on each end. Junos OS supports a non-signaled (static) configuration for Ethernet and SONET, as well as the 802.3ad standardized LACP protocol for negotiation over Ethernet links.

- Equal-Cost Multipath (ECMP) (Layer 3)

By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. There is also an option that allows multiple next-hop addresses to be installed in the forwarding table, known as per-packet load balancing.

ECMP load balancing can be:

- Across BGP paths (BGP multipath)
- Within a BGP path, across multiple LSPs

In complex Ethernet topologies, traffic imbalances occur due to increased traffic flow, and load balancing becomes challenging for some of the following reasons:

- Incorrect load balancing by aggregate next hops
- Incorrect packet hash computation
- Insufficient variance in the packet flow
- Incorrect pattern selection

As a result of traffic imbalance, the load is not well distributed causing congestion in certain links, whereas some other links are not efficiently utilized.

To overcome these challenges, Junos OS provides the following solutions for resolving the genuine traffic imbalance on aggregated Ethernet bundles (IEEE 802.3ad).

- Adaptive Load Balancing

Adaptive load balancing uses a feedback mechanism to correct a genuine traffic imbalance. To correct the imbalance weights, the bandwidth and packet stream of links are adapted to achieve efficient traffic distribution across the links in an AE bundle.

To configure adaptive load balancing, include the adaptive statement at the [edit interfaces aex aggregated-ether-options load-balance] hierarchy level.

To configure the tolerance value as a percentage, include the tolerance optional keyword at the [edit interfaces aex aggregated-ether-options load-balance adaptive] hierarchy level.

To configure adaptive load balancing based on packets per second (instead of the default bits per second setting), include the `pps` optional keyword at the `[edit interfaces aex aggregated-ether-options load-balance adaptive]` hierarchy level.

To configure the scan interval for the hash value based on the sample rate for the last two seconds, include the `scan-interval` optional keyword at the `[edit interfaces aex aggregated-ether-options load-balance adaptive]` hierarchy level.

- **Per-Packet Random Spray Load Balancing**

When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of an AE bundle are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

To configure per-packet random spray load balancing, include the `per-packet` statement at the `[edit interfaces aex aggregated-ether-options load-balance]` hierarchy level.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being used by issuing the `show interfaces aex aggregated-ether-options load-balance` command.

SEE ALSO

show interfaces (Aggregated Ethernet)

Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data

IN THIS SECTION

- [Guidelines for Configuring Stateful Load Balancing for Aggregated Ethernet Interfaces or LAG Bundles | 53](#)

When multiple flows are transmitted out of an aggregated Ethernet (`ae`) interface, the flows must be distributed across the different member links evenly to enable an effective and optimal load-balancing

behavior. To obtain a streamlined and robust method of load-balancing, the member link of the aggregated Ethernet interface bundle that is selected each time for load balancing plays a significant part. The balanced mode of link selection uses 'n' bits in a precomputed hash value if it needs to select one of 2^n (2 raised to the power of n) next-hop in the unilist. The unbalanced mode of member-link or next-hop selection uses 8 bits in a precomputed hash to select an entry in a selector table, which is randomly done with the member link IDs of the link aggregation group (LAG) or aebundle.

The term balanced versus unbalanced indicates whether a selector table is used for load balancing mechanism or not. The LAG bundle uses the unbalanced mode (selector table balancing) to balance the traffic across member links. When the traffic flows are minimal, the following problems might occur with the unbalanced mode: The link selection logic utilizes only subset bits of the precomputed hash. Regardless of the efficiency of the hashing algorithm, it is only the compressed representation of a flow. Because the inter-flow variance is very low, the resultant hashes and the subset that are computed do not provide the necessary variability to effectively utilize all the LAG member links. An excessive amount of random nature exists in the hash computation and also in the selector table. As a result, the deviation from being an optimal load-balancing technique for each child link that is selected is higher when the number of flows is lower.

The deviation per child link is defined as

$$Vi = ((Ci - (M/N))/N)$$

where

- Vi denotes the deviation for that child link 'i'.
- i denotes the child link member/index.
- Ci represents the packets transmitted for that child link 'i'.
- M signifies the total packets transmitted on that LAG bundle.
- N denotes the number of child links in that LAG.

Because of these drawbacks, for smaller number of flows, or flows with less inter-flow variance, the link utilization is skewed, and a high probability of a few child links not being utilized entirely exists.

The mechanism to record and retain states for the flows and distribute the traffic load accordingly is added. As a result, for m number of flows, they are distributed among n member links of a LAG bundle or among the unilist of next-hops in an ECMP link. This method of splitting the load among member links is called *stateful load balancing* and it uses 5-tuple information (source and destination addresses, protocol, source and destination ports). Such a method can be mapped directly to the flows, or to a precompute hash based on certain fields in the flow. As a result, the deviation observed on each child link is reduced.

This mechanism works efficiently only for minimal number of flows (less than thousands of flows, approximately). For a larger number of flows (between 1000 and 10,000 flows), we recommend that distributed Trio-based load-balancing mechanism is used.

Consider a sample scenario in which 'n' links in the LAG are identified with link IDs of 0 through n-1. A hash table or a flow table is used to record the flows as and when they show up. The hashing key is constructed using the fields that uniquely identify a flow. The result of the lookup identifies the link_id that the flow is currently using. For each packet, the flow table based on the flow identifier is examined. If a match is found, it denotes a packet that belongs to a flow that is previously processed or detected. The link ID is associated with the flow. If a match is not found, it is the first packet that belongs to the flow. The link ID is used to select the link and the flow is inserted into the flow table.

To enable per-flow load balancing based on hash values, include the per-flow statement at the [edit interfaces aeX unit *logical-unit-number* forwarding-options load-balance-stateful] hierarchy level. By default, Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. All Packet Forwarding Engine slots are assigned the same hash value by default. To configure the load-balancing algorithm to dynamically rebalance the LAG using existing parameters, include the rebalance *interval* statement at the [edit interfaces aeX unit *logical-unit-number* forwarding-options load-balance-stateful] hierarchy level. This parameter periodically load balances traffic by providing a synchronized rebalance switchover across all the ingress Packet Forwarding Engines (PFEs) over a rebalance interval. You can specify the interval as a value in the range of 1 through 1000 flows per minute. To configure the load type, include the load-type (low | medium | high) statement at the [edit interfaces aeX unit *logical-unit-number* forwarding-options load-balance-stateful] hierarchy level.

The stateful per-flow option enables the load-balancing capability on AE bundles. The rebalance option clears the load balance state at specified intervals. The load option informs the Packet Forwarding Engine regarding the appropriate memory pattern to be used. If the number of flows that flow on this aggregated Ethernet interface is less (between 1 and 100 flows), then the low keyword can be used. Similarly for relatively higher flows (between 100 and 1000 flows), the medium keyword can be used and the large keyword can be used for the maximum flows (between 1000 and 10,000 flows). The approximate number of flows for effective load-balancing for each keyword is a derivative.

The clear interfaces aeX unit *logical-unit-number* forwarding-options load-balance state command clears the load balance state at the hardware level and enables rebalancing from the cleaned up, empty state. This clear state is triggered only when you use this command. The clear interfaces aggregate forwarding-options load-balance state command clears all the aggregate Ethernet interface load balancing states and re-creates them newly.

Guidelines for Configuring Stateful Load Balancing for Aggregated Ethernet Interfaces or LAG Bundles

Keep the following points in mind while configuring stateful load-balancing for aggregated Ethernet interfaces:

- When a child link is removed or added, a new aggregate selector is selected and traffic flows onto the new selector. Because the selector is empty, flows are filled in the selector. This behavior causes redistribution of flows because the old state is lost. This is the existing behavior without enabling stateful per-flow load-balancing.
- Stateful per-flow load-balancing functions on AE interfaces if the incoming traffic reaches the MPC1E, MPC2E, MPC3E-3D, MPC5E, and MPC6E line cards. Any other type of line card does not trigger this functionality. Appropriate CLI errors are displayed if the MPCs do not support this capability.

With the ingress line card as MPC and the egress line card as MPC or DPC, this feature works properly. Stateful load-balancing is not supported if the ingress line card is a DPC and the egress line card is a DPC or an MPC.

- This capability is not supported for multicast traffic (native/flood).
- Enabling the rebalance option or clearing the load balance state can cause packet reordering for active flows because different sets of links can be selected for traffic flows.
- Although the feature performance is high, it consumes significant amount of line card memory. Approximately, 4000 logical interfaces or 16 aggregated Ethernet logical interfaces can have this feature enabled on supported MPCs. However, when the Packet Forwarding Engine hardware memory is low, depending upon the available memory, it falls back to the default load balancing mechanism. A system logging message is generated in such a situation and sent to the Routing Engine. A restriction on the number of AE interfaces that support stateful load-balancing does not exist; the limit is determined by the line cards.
- If the traffic flows become aged frequently, then the device needs to remove or refresh the load balancing states. As a result, you must configure rebalancing or run the clear command at periodic intervals for proper load-balancing. Otherwise, traffic skewing can occur. When a child link goes down or comes up, the load balancing behavior does not undergo changes on existing flows. This condition is to avoid packet reordering. New flows pick up the child link that come up. If you observe load distribution to be not very effective, you can clear the load-balancing states or use rebalancing functionality to cause an automatic clearance of the hardware states. When you configure the rebalancing facility, traffic flows can get redirected to different links, which can cause packet reordering.

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces

The mechanism to record and retain states for the flows and distribute the traffic load accordingly is added. As a result, for m number of flows, they are distributed among n member links of a LAG bundle or among the unilist of next-hops in an ECMP link. This method of splitting the load among member links is called *stateful load balancing* and it uses 5-tuple information (source and destination addresses,

protocol, source and destination ports). Such a method can be mapped directly to the flows, or to a precompute hash based on certain fields in the flow. As a result, the deviation observed on each child link is reduced.

To configure stateful load balancing on ae interface bundles:

1. Specify that you want to configure an aggregated Ethernet interface.

```
[edit]
user@R2# set interfaces aeX unit logical-unit-number
```

2. Specify that you want to configure stateful load-balancing.

```
[edit interfaces aeX unit logical-unit-number]
user@R2# edit forwarding-options load-balance-stateful
```

3. Enable the mechanism to perform an even, effective distribution of traffic flows across member links of an aggregated Ethernet interface (ae) bundle on MX Series routers with MPCs, except MPC3Es and MPC4Es.

```
[edit interfaces aeX unit logical-unit-number load-balance-stateful]
user@R2# set per-flow
```

4. Configure periodic rebalancing of traffic flows of an aggregated Ethernet bundle by clearing the load balance state at a specified interval.

```
[edit interfaces aeX unit logical-unit-number load-balance-stateful]
user@R2# set rebalance interval
```

5. Define the load-balancing type to inform the Packet Forwarding Engine regarding the appropriate memory pattern to be used for traffic flows. The approximate number of flows for effective load-balancing for each keyword is a derivative.

```
[edit interfaces aeX unit logical-unit-number load-balance-stateful]
user@R2# set load-type (low | medium | large)
```

6. Configure the address family and IP address for the ae interface.

```
[edit interfaces ae<unit logical-unit-number>]
user@R2# set family family-name address address
```

Configuring Adaptive Load Balancing

This topic describes how to configure adaptive load balancing. Adaptive load balancing maintains efficient utilization of member link bandwidth for an aggregated Ethernet (AE) bundle. Adaptive load balancing uses a feedback mechanism to correct traffic load imbalance by adjusting the bandwidth and packet streams on links within an AE bundle.

Before you begin:

- Configure a set of interfaces with a protocol family and IP address. These interfaces can make up the membership for the AE bundle.
- Create an AE bundle by configuring a set of router interfaces as aggregated Ethernet and with a specific AE group identifier.

To configure adaptive load balancing for an AE bundles:

1. Enable adaptive load balancing on the AE bundle:

```
[edit interfaces ae-x aggregated-ether-options load-balance]
user@router# set adaptive
```

2. Configure the scan interval value for adaptive load balancing on the AE bundle. The scan interval value determines the length of the traffic scan by multiplying the integer value with a 30-second time period:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set scan-interval multiplier
```

- Configure the tolerance percentage value. The tolerance value determines the allowed deviation in the traffic rates among the members of the AE bundle before the router triggers an adaptive load balancing update:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set tolerance percentage
```

- (Optional) Enable packet-per-second-based adaptive load balancing on the AE bundle:

```
[edit interfaces ae-x aggregated-ether-options load-balance adaptive]
user@router# set pps
```

SEE ALSO

adaptive

Understanding Symmetric Hashing for Load Balancing

IN THIS SECTION

- [Benefits of Symmetric Hashing | 57](#)

On devices that support this feature, symmetric hashing keeps the forward and reverse directions of a flow on the same path for features that depend on traffic symmetry. Symmetric hashing helps services maintain session context while load balancing across ECMP next hops and LAG members.

Benefits of Symmetric Hashing

- Maintains bidirectional path symmetry so stateful services see both directions of a flow on the same path.
- Improves load-balancing predictability by using the same normalized inputs for both directions of a flow.
- Reduces asymmetric routing issues that can disrupt service processing or session state.

To configure symmetric hashing, use the `symmetric-hash` configuration statement at the `[edit forwarding-options enhanced-hash-key]` hierarchy. Use the `enhanced-hash-key` statement to configure the packet fields included in hash calculations. By selecting stable, bidirectional fields and excluding unidirectional inputs, you normalize the inputs so the device derives the same hash for both directions of a flow.

You can exclude unidirectional attributes, such as the ingress interface, from the hash with `no-incoming-port` configuration statement at the `[edit forwarding-options enhanced-hash-key]` hierarchy.

Verify the current hash inputs and symmetric hashing status with the `show forwarding-options enhanced-hash-key` command.

Additional considerations:

- Dynamic load balancing (DLB) takes precedence. Symmetric hashing applies to static load balancing; if DLB is enabled, DLB overrides symmetric behavior.
- Use Layer 3 and Layer 4 fields for symmetric hashing. Layer 2 header fields are not included in symmetric hashing and you should use the `hash-mode` configuration statement to use the Layer 2 data field when required.
- Hash polarization can occur in cascaded topologies if all devices use identical hash seeds and inputs. Consider resilient hashing and seed diversity to reduce polarization when links flap or members change.

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

IN THIS SECTION

- [Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview | 58](#)
- [Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers | 59](#)
- [Configuring Symmetrical Load Balancing on Trio-Based MPCs | 62](#)
- [Example Configurations | 64](#)

Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview

MX Series routers with Aggregated Ethernet PICs support symmetrical load balancing on an 802.3ad LAG. This feature is significant when two MX Series routers are connected transparently through deep

packet inspection (DPI) devices over an LAG bundle. DPI devices keep track of flows and require information of a given flow in both forward and reverse directions. Without symmetrical load balancing on an 802.3ad LAG, the DPIs could misunderstand the flow, leading to traffic disruptions. By using this feature, a given flow of traffic (duplex) is ensured for the same devices in both directions.

Symmetrical load balancing on an 802.3ad LAG utilizes a mechanism of interchanging the source and destination addresses for a hash computation of fields, such as source address and destination address. The result of a hash computed on these fields is used to choose the link of the LAG. The hash-computation for the forward and reverse flow must be identical. This is achieved by swapping source fields with destination fields for the reverse flow. The swapped operation is referred to as *complement hash computation* or *symmetric-hash complement* and the regular (or unswapped) operation as *symmetric-hash computation* or *symmetric-hash*. The swappable fields are MAC address, IP address, and port.

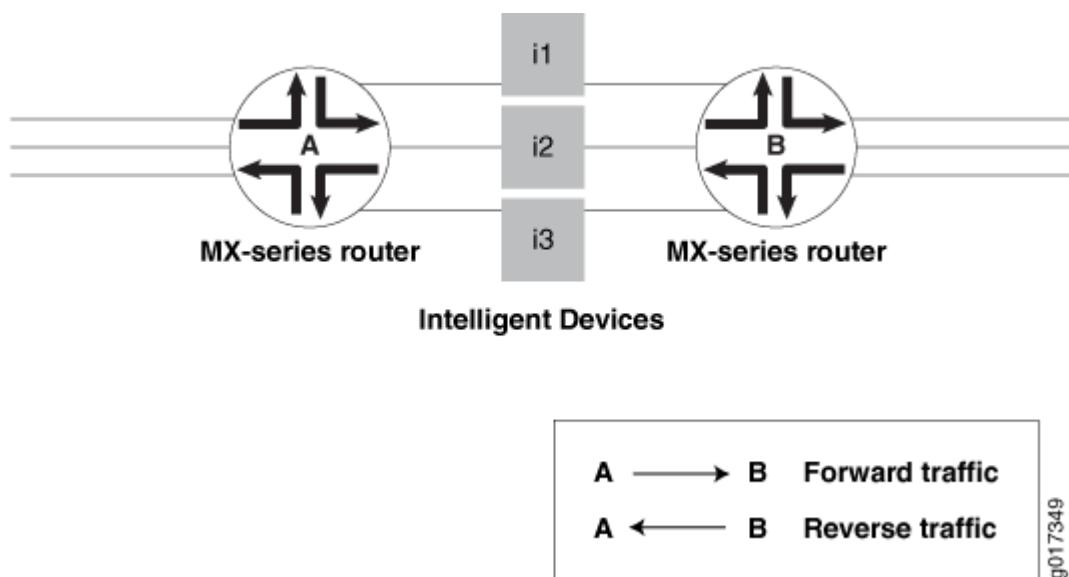
Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

You can specify whether symmetric hash or complement hash is done for load-balancing traffic. To configure symmetric hash, use the `symmetric-hash` statement at the `[edit forwarding-options hash-key family inet]` hierarchy level. To configure symmetric hash complement, use the `symmetric-hash complement` statement and option at the `[edit forwarding-options hash-key family inet]` hierarchy level.

These operations can also be performed at the PIC level by specifying a *hash key*. To configure a hash key at the PIC level, use the `symmetric-hash` or `symmetric-hash complement` statement at the `[edit chassis hash-key family inet]` and `[edit chassis hash-key family multiservice]` hierarchy levels.

Consider the example in [Figure 3 on page 59](#).

Figure 3: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers



Router A is configured with symmetric hash and Router B is configured with symmetric hash complement. Thus, for a given flow f_x , post hash computation is from Router A to Router B through i2. The reverse traffic for the same flow f_x is from Router B to Router A through the same i2 device as its hashing (done after swapping source and destination fields) and returns the same link index; since it is performed on the interchanged source and destination addresses.

However, the link chosen may or may not correspond to what was attached to the DPI. In other words, the hashing result should point to the same links that are connected, so that the traffic flows through the same DPI devices in both directions. To make sure this happens, you need to also configure the counterpart ports (ports that are connected to same DPI-iN) with the identical link index. This is done when configuring a child-link into the LAG bundle. This ensures that the link chosen for a given hash result is always the same on either router.

Note that any two links connected to each other should have the same link index and these link indices must be unique in a given bundle.



NOTE: The following restrictions apply when configuring symmetric load balancing on an 802.3ad LAG on MX Series routers:

- The Packet Forwarding Engine (PFE) can be configured to hash the traffic in either symmetric or complement mode. A single PFE complex cannot work simultaneously in both operational modes and such a configuration can yield undesirable results.
- The per-PFE setting overrides the chassis-wide setting only for the family configured. For the other families, the PFE complex still inherits the chassis-wide setting (when configured) or the default setting.
- This feature supports VPLS, INET, and bridged traffic only.
- This feature cannot work in tandem with the per-flow-hash-seed load-balancing option. It requires that all the PFE complexes configured in complementary fashion share the same seed. A change in the seed between two counterpart PFE complexes may yield undesired results.

For additional information, see the [Junos OS VPNs Library for Routing Devices](#) and the [Junos OS Administration Library for Routing Devices](#).

Example Configuration Statements

To configure 802.3ad LAG parameters at the bundle level:

```
[edit interfaces]
g(x)e-fpc/pic/port {
    gigether-options {
        802.3ad {
```

```

        bundle;
        link-index number;
    }
}
}

```

where the link-index *number* ranges from 0 through 15.

You can check the link index configured above using the `show interfaces` command:

```

[edit forwarding-options hash-key]
family inet {
    layer-3;
    layer-4;
    symmetric-hash {
        [complement;]
    }
}
family multiservice {
    source-mac;
    destination-mac;
    payload {
        ip {
            layer-3 {
                source-ip-only | destination-ip-only;
            }
            layer-4;
        }
    }
    symmetric-hash {
        [complement;]
    }
}

```

For load-balancing Layer 2 traffic based on Layer 3 fields, you can configure 802.3ad LAG parameters at a per PIC level. These configuration options are available under the chassis hierarchy as follows:

```

[edit chassis]
fpc X {
    pic Y {
        .
        .
    }
}

```

```

.
hash-key {
    family inet {
        layer-3;
        layer-4;
        symmetric-hash {
            [complement;]
        }
    }
    family multiservice {
        source-mac;
        destination-mac;
        payload {
            ip {
                layer-3 {
                    source-ip-only | destination-ip-only;
                }
                layer-4;
            }
        }
        symmetric-hash {
            [complement;]
        }
    }
}
.
.
.
}
}

```

Configuring Symmetrical Load Balancing on Trio-Based MPCs

With some configuration differences, symmetrical load-balancing over an 802.3ad link aggregation group is supported on MX Series routers with Trio-based MPCs.

To achieve symmetrical load-balancing on Trio-Based MPCs, the following needs to be done:

- Compute a Symmetrical Hash

Both routers must compute the same hash value from the flow in the forward and reverse directions. On Trio-based platforms, the calculated hash value is independent of the direction of the flow, and hence is always symmetric in nature. For this reason, no specific configuration is needed to compute a symmetric hash value on Trio-based platforms.

However, it should be noted that the fields used to configure the hash should have identical include and exclude settings on both ends of the LAG.

- **Configure Link Indexes**

To allow both routers to choose the same link using the same hash value, the links within the LAG must be configured with the same link index on both routers. This can be achieved with the `link-index` statement.

- **Enable Symmetric Load Balancing**

To configure symmetric load balancing on Trio-based MPCs, include the `symmetric` statement at the `[edit forwarding-options enhanced-hash-key]` hierarchy level. This statement is applicable to Trio-based platforms only.

The `symmetric` statement can be used with any protocol family and enables symmetric load-balancing for all aggregated Ethernet bundles on the router. The statement needs to be enabled at both ends of the LAG. This statement is disabled by default.

- **Achieve Symmetry for Bridged and Routed Traffic**

In some deployments, the LAG bundle on which symmetry is desired is traversed by Layer 2 bridged traffic in the upstream direction and by IPv4 routed traffic in the downstream direction. In such cases, the computed hash is different in each direction because the Ethernet MAC addresses are taken into account for bridged packets. To overcome this, you can exclude source and destination MAC addresses from the enhanced-hash-key computation.

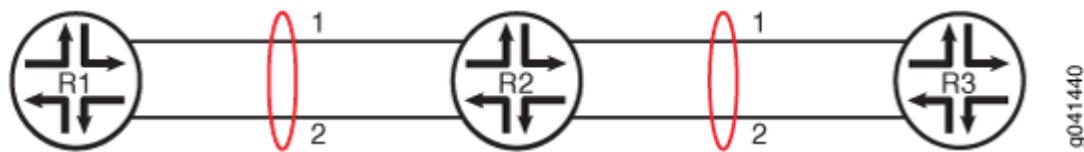
To exclude source and destination MAC addresses from the enhanced-hash-key computation, include the `no-mac-addresses` statement at the `[edit forwarding-options enhanced-hash-key family multiservice]` hierarchy level. This statement is disabled by default.

When symmetrical load balancing is enabled on Trio-based MPCs, keep in mind the following caveats:

- Traffic polarization is a phenomenon that occurs when using topologies that distribute traffic by using hashing of the same type. When routers are cascaded, traffic polarization can occur, and this can lead to unequal traffic distribution.

Traffic polarization occurs when LAGs are configured on cascaded routers. For example, in [Figure 4 on page 64](#), if a certain flow uses Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, the flow also uses Link 1 of the aggregated Ethernet bundle between Device R2 and Device R3.

Figure 4: Traffic Polarization on Cascaded Routers When Symmetrical Load Balancing is Enabled on Trio-based MPCs



This is unlike having a random link selection algorithm, where a flow might use Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, and Link 2 of the aggregated Ethernet bundle between Device R2 and Device R3.

- Symmetric load balancing is not applicable to per-prefix load-balancing where the hash is computed based on the route prefix.
- Symmetric load balancing is not applicable to MPLS or VPLS traffic, because in these scenarios the labels are not the same in both directions.

Example Configurations

IN THIS SECTION

- [Example Configurations of Chassis Wide Settings | 64](#)
- [Example Configurations of Per-Packet-Forwarding-Engine Settings | 65](#)

Example Configurations of Chassis Wide Settings

Router A

```
user@host> show configuration forwarding-options hash-key
family multiservice {
    payload {
        ip {
            layer-3;
        }
    }
    symmetric hash;
}
```

Router B

```
user@host> show configuration forwarding-options hash-key
family multiservice {
    payload {
        ip {
            layer-3;
        }
    }
    symmetric-hash {
        complement;
    }
}
```

Example Configurations of Per-Packet-Forwarding-Engine Settings**Router A**

```
user@host> show configuration chassis fpc 2 pic 2 hash-key
family multiservice {
    payload {
        ip {
            layer-3;
        }
    }
    symmetric hash;
}
```

Router B

```
user@host> show configuration chassis fpc 2 pic 3 hash-key
family multiservice {
    payload {
        ip {
            layer-3;
        }
    }
    symmetric-hash {
        complement;
    }
}
```

```
    }  
}
```

RELATED DOCUMENTATION

[Junos OS VPNs Library for Routing Devices](#)

[Junos OS Administration Library for Routing Devices](#)

Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers

Symmetrical hashing for load balancing on an 802.3ad Link Aggregation Group (LAG) is useful when two MX Series routers (for example, Router A and Router B) are connected transparently through Deep Packet Inspection (DPI) devices over a LAG bundle. The DPI devices keep track of traffic flows in both the forward and reverse directions.

If symmetrical hashing is configured, the reverse flow of traffic is also directed through the same child link on the LAG and is bound to flow through the same DPI device. This enables proper accounting on the DPI of the traffic in both the forward and reverse flows.

If symmetrical hashing is not configured, a different child link on the LAG might be chosen for the reverse flow of traffic through a different DPI device. This results in incomplete information about the forward and reverse flows of traffic on the DPI device leading to incomplete accounting of the traffic by the DPI device.

Symmetrical hashing is computed based on fields like source address and destination address. You can configure symmetrical hashing both at the chassis level and the PIC level for load balancing based on Layer 2, Layer 3, and Layer 4 data unit fields for family inet (IPv4 protocol family) and multiservice (switch or bridge) traffic. Symmetrical hashing configured at the chassis level is applicable to the entire router, and is inherited by all its PICs and Packet Forwarding Engines. Configuring PIC-level symmetrical hashing provides you more granularity at the Packet Forwarding Engine level.

For the two routers connected through the DPI devices over a LAG bundle, you can configure **symmetric-hash** on one router and **symmetric-hash complement** on the remote-end router or vice-versa.

To configure symmetrical hashing at the chassis level, include the **symmetric-hash** or the **symmetric-hash complement** statements at the [edit forwarding-options hash-key family] hierarchy level. For information about configuring symmetrical hashing at the chassis level and configuring the link index, see the [Junos OS Network Interfaces Library for Routing Devices](#) and the [Junos OS VPNs Library for Routing Devices](#).



NOTE: On MX Series DPCs, configuring symmetrical hashing at the PIC level refers to configuring symmetrical hashing at the Packet Forwarding Engine level.

To configure symmetrical hashing at the PIC level on the inbound traffic interface (where traffic enters the router), include the **symmetric-hash** or **symmetric-hash complement** statement at the [**edit chassis fpc slot-number pic pic-number hash-key**] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number hash-key]
family multiservice {
    source-mac;
    destination-mac;
    payload {
        ip {
            layer-3 (source-ip-only | destination-ip-only);
            layer-4;
        }
    }
    symmetric-hash {
        complement;
    }
}
```

```
family inet {
    layer-3;
    layer-4;
    symmetric-hash {
        complement;
    }
}
```



NOTE:

- PIC-level symmetrical hashing overrides the chassis-level symmetrical hashing configured at the [**edit chassis forwarding-options hash-key**] hierarchy level.

- Symmetrical hashing for load balancing on 802.3ad Link Aggregation Groups is currently supported for the VPLS, INET and bridged traffic only.
- Hash key configuration on a PIC or Packet Forwarding Engine can be either in the “symmetric hash” or the “symmetric hash complement” mode, but not both at the same time.

SEE ALSO

hash-key

inet

payload

symmetric-hash

Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers

IN THIS SECTION

- Configuring Symmetrical Hashing for family multiservice on Both Routers | [69](#)
- Configuring Symmetrical Hashing for family inet on Both Routers | [70](#)
- Configuring Symmetrical Hashing for family inet and family multiservice on the Two Routers | [70](#)



NOTE: These examples are applicable only to the DPCs Supported on MX240, MX480, and MX960 Routers. For the list of DPCs supported, see *DPCs Supported on MX240, MX480, and MX960 Routers* in the Related Documentation section.

The following examples show how to configure symmetrical hashing at the PIC level for load balancing on MX Series routers:

Configuring Symmetrical Hashing for family multiservice on Both Routers

On the inbound traffic interface where traffic enters Router A, include the `symmetric-hash` statement at the `[edit chassis fpc slot-number pic pic-number hash-key family multiservice]` hierarchy level:

```
[edit chassis fpc 2 pic 2 hash-key]
family multiservice {
    source-mac;
    destination-mac;
    payload {
        ip {
            layer-3;
            layer-4;
        }
    }
    symmetric-hash;
}
```

On the inbound traffic interface where traffic enters Router B, include the `symmetric-hash complement` statement at the `[edit chassis fpc slot-number pic pic-number hash-key family multiservice]` hierarchy level:

```
[edit chassis fpc 0 pic 3 hash-key]
family multiservice {
    source-mac;
    destination-mac;
    payload {
        ip {
            layer-3;
            layer-4;
        }
    }
    symmetric-hash {
        complement;
    }
}
```

Configuring Symmetrical Hashing for family inet on Both Routers

On the inbound traffic interface where traffic enters Router A, include the `symmetric-hash` statement at the `[edit chassis fpc slot-number pic pic-number hash-key family inet]` hierarchy level:

```
[edit chassis fpc 0 pic 1 hash-key]
family inet {
    layer-3;
    layer-4;
    symmetric-hash;
}
```

On the inbound traffic interface where traffic enters Router B, include the `symmetric-hash complement` statement at the `[edit chassis fpc slot-number pic pic-number hash-key family inet]` hierarchy level:

```
[edit chassis fpc 1 pic 2 hash-key]
family inet {
    layer-3;
    layer-4;
    symmetric-hash {
        complement;
    }
}
```

Configuring Symmetrical Hashing for family inet and family multiservice on the Two Routers

On the inbound traffic interface where traffic enters Router A, include the `symmetric-hash` statement at the `[edit chassis fpc slot-number pic pic-number hash-key family multiservice]` hierarchy level:

```
[edit chassis fpc 1 pic 0 hash-key]
family multiservice {
    payload {
        ip {
            layer-3;
            layer-4;
        }
    }
}
```

```
        symmetric-hash;  
    }
```

On the inbound traffic interface where traffic enters Router B, include the `symmetric-hash complement` statement at the `[edit chassis fpc slot-number pic pic-number hash-key family inet]` hierarchy level:

```
[edit chassis fpc 0 pic 3 hash-key]  
family inet {  
    layer-3;  
    layer-4;  
    symmetric-hash {  
        complement;  
    }  
}
```

SEE ALSO

| [DPCs Supported on MX240, MX480, and MX960 Routers](#)

Example: Configuring Aggregated Ethernet Load Balancing

IN THIS SECTION

- [Example: Configuring Aggregated Ethernet Load Balancing | 71](#)

Example: Configuring Aggregated Ethernet Load Balancing

IN THIS SECTION

- [Requirements | 72](#)
- [Overview | 72](#)
- [Configuration | 74](#)

- [Verification | 88](#)

This example shows how to configure aggregated Ethernet load balancing.

Requirements

This example uses the following hardware and software components:

- Three MX Series routers with MIC and MPC interfaces or three PTX Series Packet Transport Routers with PIC and FPC interfaces
- Junos OS Release 13.3 or later running on all devices

Overview

IN THIS SECTION

- [Topology | 74](#)

Load balancing is required on the forwarding plane when there are multiple paths or interfaces available to the next hop router, and it is best if the incoming traffic is load balanced across all available paths for better link utilization.

Aggregated Ethernet bundle is a typical application that uses load balancing to balance traffic flows across the member links of the bundle (IEEE 802.3ad).

Starting with Junos OS Release 13.3, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on MICs or MPCs of MX Series routers. Starting with Junos OS Release 14.1, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on PICs or FPCs of PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are:

- Adaptive—Adaptive load balancing is used in scenarios where flow-based hashing is not sufficient to achieve a uniform load distribution. This load-balancing solution implements a real-time feedback and control mechanism to monitor and manage imbalances in network load.

The adaptive load-balancing solution corrects the traffic flow imbalance by modifying the selector entries, and periodically scanning the link utilization on each member link of the AE bundle to detect any deviations. When a deviation is detected, an adjustment event is triggered and fewer flows are mapped to the affected member link. As a result, the offered bandwidth of that member link goes down. This causes a continuous feedback loop, which over a period of time ensures that the same amount of byte rate is offered to all the member links, thus providing efficient traffic distribution across each member link in the AE bundle.

To configure adaptive load balancing, include the `adaptive` statement at the `[edit interfaces aex aggregated-ether-options load-balance]` hierarchy level.



NOTE: Enabling adaptive load balancing may cause packet reordering once every rebalance interval.

The `pps` option enables load balancing based on the packets-per-second rate. The default setting is bits-per-second load balancing.

The `scan-interval` value configures the length of time for scanning as a multiple of 30 seconds.

The `tolerance` value is the limit to the variance in the packet traffic flow to the aggregated Ethernet links in the bundle. You can specify a maximum of 100-percent variance. When the `tolerance` attribute is not configured, a default value of 20 percent is enabled for adaptive load balancing. A smaller tolerance value balances better bandwidth, but takes a longer convergence time.

- **Per-packet random spray**—When the adaptive load-balancing solution fails, per-packet random spray acts as a last resort. The per-packet random spray load-balancing solution helps to address traffic imbalance by randomly spraying the packets to the aggregate next hops. This ensures that all the member links of the AE bundle are equally loaded, resulting in packet reordering.

In addition, per-packet random spray identifies the ingress Packet Forwarding Engine that caused the traffic imbalance and eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

To configure per-packet random spray load balancing, include the `per-packet` statement at the `[edit interfaces aex aggregated-ether-options load-balance]` hierarchy level.



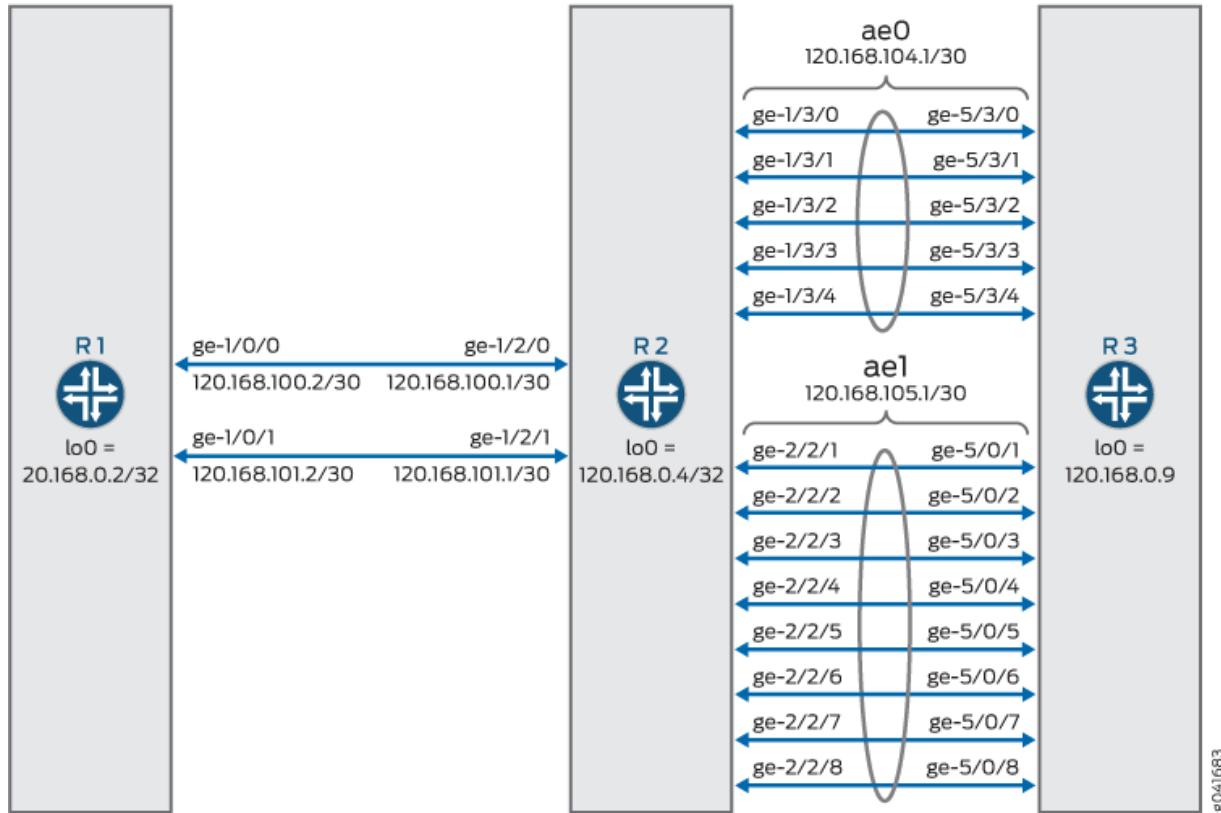
NOTE: The Per-Packet option for load balancing is not supported on the PTX Series Packet Transport Routers.

The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured one. You can verify the load-balancing solution being implemented by issuing the `show interfaces aex aggregated-ether-options load-balance` command.

Topology

In this topology, two aggregated Ethernet bundles - ae0 and ae1 - are configured on the links between the R2 and R3 routers.

Figure 5: Aggregated Ethernet Load Balancing



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 75](#)
- [Configuring Adaptive Load Balancing | 80](#)
- [Results | 83](#)

CLI Quick Configuration

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

R1

```
set chassis aggregated-devices ethernet device-count 12
set interfaces xe-0/0/0 unit 0 family inet address 120.168.1.1/30
set interfaces xe-0/0/0 unit 0 family iso
set interfaces xe-0/0/0 unit 0 family mpls
set interfaces xe-0/0/1 unit 0 family inet address 120.168.2.1/30
set interfaces xe-0/0/1 unit 0 family iso
set interfaces xe-0/0/1 unit 0 family mpls
set interfaces ge-1/0/0 unit 0 family inet address 120.168.100.2/30
set interfaces ge-1/0/0 unit 0 family iso
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces ge-1/0/1 unit 0 family inet address 120.168.101.2/30
set interfaces ge-1/0/1 unit 0 family iso
set interfaces ge-1/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0002.00
set routing-options router-id 120.168.0.2
set routing-options autonomous-system 55
set protocols rsvp interface ge-1/0/0.0
set protocols rsvp interface ge-1/0/1.0
set protocols mpls label-switched-path videl-to-sweets to 120.168.0.9
set protocols mpls label-switched-path v-2-s-601 to 60.0.1.0
set protocols mpls label-switched-path v-2-s-601 primary v-2-s-601-primary hop-limit 5
set protocols mpls label-switched-path v-2-s-602 to 60.0.2.0
set protocols mpls label-switched-path v-2-s-602 primary v-2-s-602-primary hop-limit 5
set protocols mpls label-switched-path v-2-s-603 to 60.0.3.0
set protocols mpls label-switched-path v-2-s-604 to 60.0.4.0
set protocols mpls path v-2-s-601-primary 120.168.100.1 strict
set protocols mpls path v-2-s-601-primary 120.168.104.2 strict
set protocols mpls path v-2-s-602-primary 120.168.101.1 strict
set protocols mpls path v-2-s-602-primary 120.168.105.2 strict
set protocols mpls interface ge-1/0/0.0
set protocols mpls interface ge-1/0/1.0
set protocols bgp group pe-routers type internal
set protocols bgp group pe-routers local-address 120.168.0.2
```

```

set protocols bgp group pe-routers family inet unicast
set protocols bgp group pe-routers family inet-vpn unicast
set protocols bgp group pe-routers neighbor 120.168.0.9
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ge-1/0/0.0
set protocols isis interface ge-1/0/1.0
set protocols isis interface lo0.0
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement vpn-m5-export term 1 from protocol bgp
set policy-options policy-statement vpn-m5-export term 1 from protocol direct
set policy-options policy-statement vpn-m5-export term 1 then community add vpn-m5-target
set policy-options policy-statement vpn-m5-export term 1 then accept
set policy-options policy-statement vpn-m5-export term 2 then reject
set policy-options policy-statement vpn-m5-import term 1 from protocol bgp
set policy-options policy-statement vpn-m5-import term 1 from community vpn-m5-target
set policy-options policy-statement vpn-m5-import term 1 then accept
set policy-options policy-statement vpn-m5-import term 2 then reject
set policy-options community vpn-m5-target members target:55:100
set routing-instances vpn-m5 instance-type vrf
set routing-instances vpn-m5 interface xe-0/0/0.0
set routing-instances vpn-m5 interface xe-0/0/1.0
set routing-instances vpn-m5 route-distinguisher 120.168.0.2:1
set routing-instances vpn-m5 vrf-import vpn-m5-import
set routing-instances vpn-m5 vrf-export vpn-m5-export
set routing-instances vpn-m5 protocols bgp group ce type external
set routing-instances vpn-m5 protocols bgp group ce peer-as 100
set routing-instances vpn-m5 protocols bgp group ce as-override
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.1.2
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.2.2
set routing-instances vpn-m5 protocols ospf domain-id 1.0.0.0
set routing-instances vpn-m5 protocols ospf export vpn-m5-import
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-0/0/1.0
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-0/0/0.0

```

R2

```

set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/2/0 unit 0 family inet address 120.168.100.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 120.168.101.1/30

```

```
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-1/3/0 gigether-options 802.3ad ae0
set interfaces ge-1/3/1 gigether-options 802.3ad ae0
set interfaces ge-1/3/2 gigether-options 802.3ad ae0
set interfaces ge-1/3/3 gigether-options 802.3ad ae0
set interfaces ge-1/3/4 gigether-options 802.3ad ae0
set interfaces ge-2/2/1 gigether-options 802.3ad ae1
set interfaces ge-2/2/2 gigether-options 802.3ad ae1
set interfaces ge-2/2/3 gigether-options 802.3ad ae1
set interfaces ge-2/2/4 gigether-options 802.3ad ae1
set interfaces ge-2/2/5 gigether-options 802.3ad ae1
set interfaces ge-2/2/6 gigether-options 802.3ad ae1
set interfaces ge-2/2/7 gigether-options 802.3ad ae1
set interfaces ge-2/2/8 gigether-options 802.3ad ae1
set interfaces ae0 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.1/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet address 120.168.105.1/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0004.00
set accounting-options selective-aggregate-interface-stats disable
set protocols rsvp interface ge-1/2/0.0
set protocols rsvp interface ge-1/2/1.0
set protocols rsvp interface ae0.0
set protocols rsvp interface ae1.0
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ge-1/2/0.0
set protocols isis interface ge-1/2/1.0
set protocols isis interface ae0.0
```

```
set protocols isis interface ae1.0
set protocols isis interface lo0.0
```

R3

```
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-4/0/0 unit 0 family inet address 120.168.9.1/30
set interfaces xe-4/0/0 unit 0 family mpls
set interfaces xe-4/0/1 unit 0 family inet address 120.168.10.1/30
set interfaces xe-4/0/1 unit 0 family mpls
set interfaces ge-5/0/1 gigether-options 802.3ad ae1
set interfaces ge-5/0/2 gigether-options 802.3ad ae1
set interfaces ge-5/0/3 gigether-options 802.3ad ae1
set interfaces ge-5/0/4 gigether-options 802.3ad ae1
set interfaces ge-5/0/5 gigether-options 802.3ad ae1
set interfaces ge-5/0/6 gigether-options 802.3ad ae1
set interfaces ge-5/0/7 gigether-options 802.3ad ae1
set interfaces ge-5/0/8 gigether-options 802.3ad ae1
set interfaces ge-5/3/0 gigether-options 802.3ad ae0
set interfaces ge-5/3/1 gigether-options 802.3ad ae0
set interfaces ge-5/3/2 gigether-options 802.3ad ae0
set interfaces ge-5/3/3 gigether-options 802.3ad ae0
set interfaces ge-5/3/4 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.2/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet address 120.168.105.2/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.9/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0009.00
set routing-options router-id 120.168.0.9
set routing-options autonomous-system 55
set protocols rsvp interface xe-4/0/0.0
set protocols rsvp interface xe-4/0/1.0
set protocols rsvp interface ae0.0
set protocols rsvp interface ae1.0
set protocols mpls label-switched-path to-videl to 120.168.0.2
```

```
set protocols mpls interface xe-4/0/0.0
set protocols mpls interface xe-4/0/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols bgp group pe-routers type internal
set protocols bgp group pe-routers local-address 120.168.0.9
set protocols bgp group pe-routers family inet unicast
set protocols bgp group pe-routers family inet-vpn unicast
set protocols bgp group pe-routers neighbor 120.168.0.2
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ae0.0
set protocols isis interface ae1.0
set protocols isis interface lo0.0
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement vpn-m5-export term 1 from protocol bgp
set policy-options policy-statement vpn-m5-export term 1 from protocol direct
set policy-options policy-statement vpn-m5-export term 1 then community add vpn-m5-target
set policy-options policy-statement vpn-m5-export term 1 then accept
set policy-options policy-statement vpn-m5-export term 2 then reject
set policy-options policy-statement vpn-m5-import term 1 from protocol bgp
set policy-options policy-statement vpn-m5-import term 1 from protocol direct
set policy-options policy-statement vpn-m5-import term 1 from community vpn-m5-target
set policy-options policy-statement vpn-m5-import term 1 then accept
set policy-options policy-statement vpn-m5-import term 2 then reject
set policy-options community vpn-m5-target members target:55:100
set routing-instances vpn-m5 instance-type vrf
set routing-instances vpn-m5 interface xe-4/0/0.0
set routing-instances vpn-m5 interface xe-4/0/1.0
set routing-instances vpn-m5 route-distinguisher 120.168.0.9:1
set routing-instances vpn-m5 vrf-import vpn-m5-import
set routing-instances vpn-m5 vrf-export vpn-m5-export
set routing-instances vpn-m5 protocols bgp group ce type external
set routing-instances vpn-m5 protocols bgp group ce peer-as 100
set routing-instances vpn-m5 protocols bgp group ce as-override
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.9.2
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.10.2
set routing-instances vpn-m5 protocols ospf domain-id 1.0.0.0
set routing-instances vpn-m5 protocols ospf export vpn-m5-import
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/0.0
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/1.0
```

Configuring Adaptive Load Balancing

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the R2 router:



NOTE: Repeat this procedure for the other routers, after modifying the appropriate interface names, addresses, and any other parameters for each router.

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

```
[edit chassis]
user@R2# set aggregated-devices ethernet device-count 5
```

2. Configure the Gigabit Ethernet interface link connecting R2 to R1.

```
[edit interfaces]
user@R2# set ge-1/2/0 unit 0 family inet address 120.168.100.1/30
user@R2# set ge-1/2/0 unit 0 family iso
user@R2# set ge-1/2/0 unit 0 family mpls
user@R2# set ge-1/2/1 unit 0 family inet address 120.168.101.1/30
user@R2# set ge-1/2/1 unit 0 family iso
user@R2# set ge-1/2/1 unit 0 family mpls
user@R2# set lo0 unit 0 family inet address 120.168.0.4/32
user@R2# set lo0 unit 0 family iso address 49.0001.1201.6800.0004.00
```

3. Configure the five member links of the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ge-1/3/0 gigether-options 802.3ad ae0
user@R2# set ge-1/3/1 gigether-options 802.3ad ae0
user@R2# set ge-1/3/2 gigether-options 802.3ad ae0
user@R2# set ge-1/3/3 gigether-options 802.3ad ae0
user@R2# set ge-1/3/4 gigether-options 802.3ad ae0
```

4. Configure the eight member links of the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ge-2/2/1 gigether-options 802.3ad ae1
user@R2# set ge-2/2/2 gigether-options 802.3ad ae1
user@R2# set ge-2/2/3 gigether-options 802.3ad ae1
user@R2# set ge-2/2/4 gigether-options 802.3ad ae1
user@R2# set ge-2/2/5 gigether-options 802.3ad ae1
user@R2# set ge-2/2/6 gigether-options 802.3ad ae1
user@R2# set ge-2/2/7 gigether-options 802.3ad ae1
user@R2# set ge-2/2/8 gigether-options 802.3ad ae1
```

5. Enable aggregate Ethernet load balancing on ae0 of R2.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options load-balance adaptive tolerance 10
```

6. Configure the link speed for the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options link-speed 1g
```

7. Configure LACP on the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 aggregated-ether-options lacp active
```

8. Configure the interface parameters for the ae0 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae0 unit 0 family inet address 120.168.104.1/30
user@R2# set ae0 unit 0 family iso
user@R2# set ae0 unit 0 family mpls
```

9. Enable aggregate Ethernet load balancing on ae1 of R2.

```
[edit interfaces]
user@R2# set ae1 aggregated-ether-options load-balance adaptive tolerance 10
```

10. Configure the link speed for the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae1 aggregated-ether-options link-speed 1g
```

11. Configure LACP on the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae1 aggregated-ether-options lacp active
```

12. Configure the interface parameters for the ae1 aggregated Ethernet bundle.

```
[edit interfaces]
user@R2# set ae1 unit 0 family inet address 120.168.105.1/30
user@R2# set ae1 unit 0 family iso
user@R2# set ae1 unit 0 family mpls
```

13. Disable selective aggregate Ethernet statistics.

```
[edit accounting-options]
user@R2# set selective-aggregate-interface-stats disable
```

14. Configure RSVP on all the interfaces of R2 and on the AE bundles.

```
[edit protocols]
user@R2# set rsvp interface ge-1/2/0.0
user@R2# set rsvp interface ge-1/2/1.0
user@R2# set rsvp interface ae0.0
user@R2# set rsvp interface ae1.0
```

15. Configure MPLS on all the interfaces of R2 and on the AE bundles.

```
[edit protocols]
user@R2# set mpls interface ge-1/2/0.0
user@R2# set mpls interface ge-1/2/1.0
user@R2# set mpls interface ae0.0
user@R2# set mpls interface ae1.0
```

16. Configure IS-IS on all the interfaces of R2 and on the AE bundles.

```
[edit protocols]
user@R2# set isis traffic-engineering family inet shortcuts
user@R2# set isis level 1 disable
user@R2# set isis interface ge-1/2/0.0
user@R2# set isis interface ge-1/2/1.0
user@R2# set isis interface ae0.0
user@R2# set isis interface ae1.0
user@R2# set isis interface lo0.0
```

Results

From configuration mode, confirm your configuration by entering the show chassis, show interfaces, show accounting-options, and show protocols commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show chassis
aggregated-devices {
    ethernet {
        device-count 5;
    }
}
```

```
user@R2# show interfaces
ge-1/2/0 {
    unit 0 {
        family inet {
            address 120.168.100.1/30;
        }
    }
}
```

```
    family iso;
    family mpls;
}
}
ge-1/2/1 {
    unit 0 {
        family inet {
            address 120.168.101.1/30;
        }
        family iso;
        family mpls;
    }
}
ge-1/3/0 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/1 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/2 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/3 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/3/4 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/2/1 {
    gigether-options {
        802.3ad ae1;
    }
}
```

```
ge-2/2/2 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/3 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/4 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/5 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/6 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/7 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/2/8 {
    gigether-options {
        802.3ad ae1;
    }
}
ae0 {
    aggregated-ether-options {
        load-balance {
            adaptive tolerance 10;
        }
        link-speed 1g;
        lacp {
            active;
        }
    }
}
```

```
        }
    }
    unit 0 {
        family inet {
            address 120.168.104.1/30;
        }
        family iso;
        family mpls;
    }
}
ae1 {
    aggregated-ether-options {
        load-balance {
            adaptive tolerance 10;
        }
        link-speed 1g;
        lacp {
            active;
        }
    }
    unit 0 {
        family inet {
            address 120.168.105.1/30;
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 120.168.0.4/32;
        }
        family iso {
            address 49.0001.1201.6800.0004.00;
        }
    }
}
```

```
    }  
}
```

```
user@R2# show accounting-options  
selective-aggregate-interface-stats disable;
```

```
user@R2# show protocols  
rsvp {  
    interface ge-1/2/0.0;  
    interface ge-1/2/1.0;  
    interface ae0.0;  
    interface ae1.0;  
}  
mpls {  
    interface ge-1/2/0.0;  
    interface ge-1/2/1.0;  
    interface ae0.0;  
    interface ae1.0;  
}  
isis {  
    traffic-engineering {  
        family inet {  
            shortcuts;  
        }  
    }  
    level 1 disable;  
    interface ge-1/2/0.0;  
    interface ge-1/2/1.0;  
    interface ae0.0;  
    interface ae1.0;  
    interface lo0.0;  
}
```

Verification

IN THIS SECTION

- [Verifying Adaptive Load Balancing on ae0 | 88](#)

Confirm that the configuration is working properly.

Verifying Adaptive Load Balancing on ae0

Purpose

Verify that packets received on the ae0 aggregated Ethernet bundle are load-balanced among the five member links.

Action

From operational mode, run the `show interfaces ae0 extensive` command.

```
user@R2> show interfaces ae0 extensive
Logical interface ae0.0 (Index 325) (SNMP ifIndex 917) (Generation 134)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  Statistics          Packets          pps          Bytes          bps
  Bundle:
    Input :      848761          9      81247024          7616
    Output: 166067308909  3503173 126900990064983  21423804256
  Adaptive Statistics:
    Adaptive Adjusts:      264
    Adaptive Scans :      27682
    Adaptive Updates:      10
  Link:
    ge-1/3/0.0
      Input :      290888          5      29454436          3072
      Output: 33183442699  704569 25358563587277  4306031760
    ge-1/3/1.0
      Input :      162703          1      14806325          992
      Output: 33248375409  705446 25406995966732  4315342152
    ge-1/3/2.0
```

Input :	127448	1	12130566	992
Output:	33184552729	697572	25354827700261	4267192376
ge-1/3/3.0				
Input :	121044	1	11481262	1280
Output:	33245875402	697716	25405953405192	4265750584
ge-1/3/4.0				
Input :	146678	1	13374435	1280
Output:	33205071207	697870	25374651121458	4269487384

Meaning

The member links of the ae0 aggregated Ethernet bundle are fully utilized with adaptive load balancing.

Platform-Specific Aggregated Ethernet Load Balancing Behavior

IN THIS SECTION

- Platform-Specific Aggregated Ethernet Load Balancing Behavior | 89

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Use the following table to review platform-specific behaviors for your platform.

Platform-Specific Aggregated Ethernet Load Balancing Behavior

Platform	Difference
ACX Series	<ul style="list-style-type: none"> On ACX7000 Series of devices, ae member interfaces do not load balance egress traffic. On ACX7000 Series of devices, you must configure the set forwarding-options hash-key statement to use all available member interfaces for load balancing.
EX Series	<ul style="list-style-type: none"> You can configure up 480 LAG bundles on EX9200 switches.

(Continued)

Platform	Difference
MX Series	<ul style="list-style-type: none"> • You can configure up 480 LAG bundles on MX Series routers that support this feature. • You can perform uniform load balancing and rebalancing on MX Series routers with MPCs that support this feature. Rebalancing is not supported when load-balancing is skewed or distorted owing to a change in the number of flows.
PTX Series	<ul style="list-style-type: none"> • Adaptive load balancing is not supported on PTX Series devices if the VLAN ID is configured on the aggregated Ethernet interface. • The pps and scan-interval optional keywords are supported on PTX Series Packet Transport Routers only.
QFX Series	<ul style="list-style-type: none"> • Adaptive load balancing is not supported on QFX10000 switches if the VLAN ID is configured on the aggregated Ethernet interface.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
14.1	Starting with Junos OS Release 14.1, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on PICs or FPCs of PTX Series Packet Transport Routers.
13.3	Starting with Junos OS Release 13.3, aggregated Ethernet load balancing is enhanced to provide two solutions for resolving genuine traffic imbalance on aggregated Ethernet bundles on MICs or MPCs of MX Series routers.
10.1	Starting with Junos OS Release 10.1, you can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the payload statement.

Performance Monitoring on Aggregated Ethernet Interfaces

SUMMARY

Learn about performance monitoring features on aggregated Ethernet Interfaces. Refer the guidelines for configuring performance monitoring features before you configure performance monitoring.

IN THIS SECTION

- [ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces | 91](#)
- [Guidelines to Configure Performance Monitoring on Aggregated Ethernet Interfaces | 92](#)

ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces

Configure ITU-T Y.1731 standard-compliant Ethernet loss measurement (ETH-LM), Ethernet synthetic loss measurement (ETH-SLM), and Ethernet delay measurement (ETH- DM) capabilities on aggregated Ethernet (ae) interfaces. These ITU-T Y.1731 OAM services or performance monitoring techniques can be measured by on-demand mode (triggered through the CLI) or by proactive mode (triggered by the iterator application).

Connectivity fault management (CFM) sessions established on the AE interfaces can be distributed to the Packet Forwarding Engine, apart from being handled on the Routing engine. This capability to distribute CFM sessions is useful in both scaled topologies and graceful Routing Engine switchover (GRES) for CFM sessions.

To enable the distribution of CFM sessions and to operate in centralized mode, include the ppm delegate-processing statement at the [edit routing-options ppm] hierarchy level. The mechanism that enables distribution of CFM sessions over AE interfaces provides the underlying infrastructure to support PM over AE interfaces. In addition, periodic packet management (PPM) handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With PPM processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run performance monitoring processes on the Packet Forwarding Engine.

For Ethernet delay measurement, hardware-assisted timestamping is supported on AE interfaces, similar to the support that exists on non-AE interfaces. Only hardware-based timestamping is supported

because it is performed in the received path of the protocol data unit (PDU) packets, whereas software-based timestamping needs to be performed on the transmitted path and is not supported.

Before you start an ETH-DM, ETH-LM, or ETH-SLM measurement sessions across an aggregated Ethernet service, you must configure two routers to support these measurement sessions. On each router, configure two physical or logical AE interfaces connected by a VLAN by including the `interface ae-fpc/pic/port unit logical-unit-number vlan-id vlan-id` statement at the [edit interfaces] hierarchy level and on each router, attach the peer MEPs to the interfaces by including the `mep mep-id interface interface-name (protect | working)` statement at the [edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association ma-name] hierarchy level.

Ethernet loss measurement over an aggregated Ethernet (ae) interface is not supported when the enhanced LAG functionality is enabled on a router. The enhanced LAG capability is enabled by default when you configure enhanced IP services mode by including the `network-services enhanced-ip` statement at the [edit chassis] hierarchy level. For Ethernet loss measurement to work properly, you must disable the enhanced LAG functionality by entering the `set chassis aggregated-devices disable-lag-enhanced` statement.

Performance monitoring for connectivity fault management is not supported when the network-to-network (NNI) or egress interface is an aggregated Ethernet interface with member links on DPCs.

Guidelines to Configure Performance Monitoring on Aggregated Ethernet Interfaces

Keep the following points in mind while you configure ETH-LM, ETH-SLM, and ETH-DM capabilities on aggregated Ethernet (ae-) interfaces:

- The scaling limits and performance considerations for distributed periodic packet management (PPM) sessions. The scaling limits for distributed PPM sessions over aggregated Ethernet (AE) interfaces are identical to the maximum supported numbers for continuity check messages (CCM) over AE interfaces.
- SLA iterators always coexist with CCM sessions. Therefore, while configuring a scaled environment, you must account for CCM sessions should be accounted along with SLA iterators. The following table describes the maximum number of distributed PM sessions you can configure for different CCM intervals per line card and per router (system-wide value).
- A mixed operation of distributed and centralized modes for performance monitoring (PM) sessions is not supported on AE interfaces, if the interfaces that form the aggregated Ethernet bundle are in mixed mode.
- The limitations for performance monitoring (PM) capabilities for non-AE interfaces apply equally well for AE interfaces. For example, flapping of sessions resets the PM statistics.

- The limitations that exist with distributed PPM sessions are valid for performance monitoring capabilities over AE interfaces because measurements are always performed on CCM sessions.
- For ETH-LM over AE interfaces in an active-standby setup, if active and standby line cards are swapped, then the measurements during this window are ignored. Whenever the link failover from the active interface to the standby interface happens, the counters are reset.
- For ETH-DM over AE interfaces, the additional time that is taken for packet transmission (packets are redirected to anchor in the received [Rx] direction and to the active child FPC in the transmitted [Tx] direction) is computed in the delay measurement.

Periodic Packet Management on Ethernet Interfaces

SUMMARY

Learn about Periodic packet management (PPM) and how to configure PPM.

IN THIS SECTION

- [Periodic Packet Management Overview | 93](#)
- [Configure Periodic Packet Management | 94](#)

Periodic Packet Management Overview

Periodic packet management (PPM) for routers is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes on the router can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the processes that control the Link Aggregation Control Protocol (LACP) and Bidirectional Forwarding Detection (BFD) protocols, and also for receiving packets on behalf of these client processes. To enable PPM to send and receive packets on their behalf, the clients establish adjacencies with PPM. When packets are not received from the client, the adjacency is marked as down and the client is informed.

PPM operates in two modes:

- **Centralized** When PPM is operating in centralized mode, it runs on the Routing Engine only.

Distributed When PPM is operating in distributed mode, it runs on the Packet Forwarding Engine. Currently, Bidirectional Forwarding Detection (BFD), Link Aggregation Control Protocol (LACP), Link Fault Management (LFM), Connectivity Fault Management (CFM), and Virtual Router Redundancy Protocol (VRRP) operate in distributed mode, by default.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only. You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.

We recommend that, generally, you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

Configure Periodic Packet Management

IN THIS SECTION

- [Enable Centralized Periodic Packet Management | 94](#)

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

Enable Centralized Periodic Packet Management

When you enable centralized periodic packet management, the `ppm` process runs on the routing engine only. When you enable centralized periodic packet management, you have disabled distributed PPM. You can enable centralized periodic packet management for troubleshooting to identify if the protocol is having issues while running on distributed mode. If you do not face the issue while, the protocol is running on centralized mode, you can narrow down the issue and identify if the issue is because of Packet Forwarding Engine failure.

To enable centralized periodic packet management:

- From configuration mode, enable centralized periodic packet management by specifying the `no-delegate-processing` statement at the `[edit]` hierarchy level.

```
[edit]
user@host# set routing-options ppm no-delegate-processing
```

- Commit the configuration by using the `commit` statement.

```
[edit]
user@host# commit
```

- Clear the current active protocol session on the device by using the `clear protocol-name session` command. For example, to clear the BFD session, use the following command.

```
[edit]
user@host# run clear bfd session
```

- Verify the periodic packet management mode by using the `show routing-options ppm` command.

```
user@host# show routing-options
ppm;
Protocol: BFD, Hold time: 900, IFL-index: 359
Distributed: FALSE
BFD discriminator: 17, BFD routing table index: 0
```

In the output, the distributed field is false and so `ppm` is centralized.

Ethernet Link Aggregation

IN THIS SECTION

- Load Balancing | [97](#)
- LACP Monitoring | [98](#)
- Understanding the Algorithm Used to Hash LAG Bundle | [99](#)

Learn about Ethernet link aggregation, also known as link aggregation group (LAG), port trunking, or port bonding, is a technology that combines multiple Ethernet links into a single logical link.

Ethernet link aggregation is mechanism for increasing the bandwidth linearly and improving the resiliency of Ethernet links by bundling or combining multiple full-duplex same-speed point-to-point Ethernet links into a single virtual link. The virtual link interface is referred to as link aggregation group (LAG) or aggregated Ethernet (AE) interface. The LAG balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

To configure aggregated Ethernet interface:

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

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

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

```
[edit interfaces]
user@host# set ae0 aggregated-ether-options minimum-links number (1 - 8)
```

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

```
[edit interfaces]
user@host# set ae0 aggregated-ether-options link-speed speed (10g / 1g / 100m)
```

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

```
[edit interfaces]
user@host# set ge-1/0/0 gigether-options 802.3ad ae0
user@host# set ge-1/0/1 gigether-options 802.3ad ae0
```

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

```
[edit interfaces]
user@host# set ae0 unit 0 family inet address ip-address
```

The above procedure creates an AE interface and they would be up and ready for running the services defined on AE logical interfaces.

AE interfaces can be VLAN-tagged or untagged. You can configure flexible-vlan-tagging, native-vlan-id, and dual-tagging on AE interfaces.

Whenever there is a configuration change (AE interface to Gigabit Ethernet interfaces or vice versa), you need to remove the existing configuration, perform a commit, then add the new configuration and again commit the configuration.

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

This step changes the interface state to down and removes the configuration statements related to aex.

```
[edit]
user@host#delete interfaces aex
```

2. Delete the interface from the device count.

```
[edit]
user@host#delete chassis aggregated-devices ethernet device-count
```

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 both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled.

Load Balancing

JUNOS load-balances traffic across member links in an AE bundle based on the Layer 3 information in the packet. You can globally configure what fields are used for load-balancing for inet and MPLS.

On the routers, the inet family knobs are available at PIC level. You can configure inet family Layer 3 and Layer 4 fields to be used for load-balancing. For bridge family, Layer 2, layer 3 and Layer 4 fields to be used for load-balancing.

Routers also support load balancing across the member links using Layer 2 source MAC addresses, destination MAC addresses, or both. This can be configured at the [edit forwarding-options hash-key family

multiservice] hierarchy level. Layer 2 source MAC addresses and destination MAC addresses are used as hash-keys for load balancing.

```
[edit]
forwarding-options {
    hash-key {
        family multiservice {
            destination-mac;
            source-mac;
        }
    }
}
```

- For IP Layer 2 packets, only IP fields are used for load balancing across member links. Source MAC address and destination MAC address are not be used for load balancing.
- For non-IP Layer 2 packets, either Source MAC address or destination MAC address is used as hash-keys for load balancing.
- If you want to hash based on layer 2 fields, then you need to configure `multiservice`.
- If you want to hash based on layer 3 and layer 4 fields, then you need to configure family (inet | inet6)

LACP Monitoring

LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP is designed to achieve the following:

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

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

LACP monitoring can be either distributed or centralized. The default is distributed and it can be overriden by configuring the centralized knob under LACP protocols. LACP exchanges are made

between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

By default, LACP does not initiate a LACP PDU exchange. LACP packets can be configured to exchange LACP PDUs at a rate of 1 packet per second, or a slower rate of 1 packet for 30 seconds.

The LACP mode can be active or passive. 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. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is turned off on aggregated Ethernet interfaces. If LACP is configured, it is in passive mode by default. To initiate transmission of LACP packets and response to LACP packets, you must configure LACP in active mode.

To enable LACP active mode, include the `lacp` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level, and specify the **active** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
    active;
}
```

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

To restore the default behavior, include the `lacp` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level, and specify the **passive** option:

```
[edit interfaces interface-name aggregated-ether-options]
lacp {
    passive;
}
```

Understanding the Algorithm Used to Hash LAG Bundle

Routers use a hashing algorithm to determine how to forward traffic over a link aggregation group (LAG) bundle.

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.

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle.

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.

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 router. The hashing algorithm recognizes the following EtherTypes:

- IPv4
- MPLS

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.

[Table 1 on page 100](#) describes the fields used for hashing by Layer 2 services. The table explains the default behavior and the configurable fields based on the type of traffic received on the Layer 2 service.

Table 1: Hashing Behavior for Pseudowire (Layer 2 Circuit) and Bridging Services

Traffic Type	Default Hash Fields	Configurable Fields (Hash keys)
Layer 2	None	Source MAC Address Destination MAC Source MAC and Destination MAC

Table 1: Hashing Behavior for Pseudowire (Layer 2 Circuit) and Bridging Services (*Continued*)

Traffic Type	Default Hash Fields	Configurable Fields (Hash keys)
IP	Source IP and Destination IP	Source MAC Address Destination MAC Source MAC and Destination MAC
MPLS	MPLS label 1 and MPLS label 2	Source MAC Address Destination MAC Source MAC and Destination MAC

[Table 2 on page 101](#) describes the fields used for hashing by Layer 3 services. The table explains the default behavior and the configurable fields based on the type of traffic received on the Layer 3 service.

Table 2: Hashing Behavior for IP Services

Traffic Type	Default Hash Fields	Configurable Fields (Hash keys)
IP	Source IP and Destination IP	Layer 3 (Source IP and/or destination IP) Layer 4 (UDP/TCP source port andr UDP/TCP destination port)

RELATED DOCUMENTATION

[Controlling Network Access Using Traffic Policing Overview](#)

3

PART

Gigabit Ethernet Interfaces

- Gigabit Ethernet PICs | 103
- Gigabit Ethernet MICs | 124
- Smart SFPs To Transport Legacy Network Traffic over Packet Switched Networks | 127
- Layer 2 Overhead Attribute in Interface Statistics | 142
- Gigabit Ethernet Policers | 150
- Gigabit Ethernet Autonegotiation | 164

Gigabit Ethernet PICs

SUMMARY

Learn about Gigabit Ethernet (GbE) PICs that are used to connect computers and servers in local area networks (LANs). How to configure GbE PICs to perform traffic policing without the need to configure a firewall filter.

IN THIS SECTION

- [P2-10G-40G-QSFPP PIC Overview | 103](#)
- [Configure the P2-10G-40G-QSFPP PIC | 114](#)
- [Example: Configure the P2-10G-40G-QSFPP PIC | 118](#)
- [Framing Overview | 122](#)
- [Understanding WAN Framing | 123](#)
- [Configure Ethernet Framing | 123](#)

Gigabit Ethernet PIC supports flexible encapsulation and MAC accounting. Gigabit Ethernet connects computers and servers in local networks. The data transfer speed and cabling improvements have recommended many enterprises to replace Fast Ethernet with Gigabit Ethernet for wired local networks.

P2-10G-40G-QSFPP PIC Overview

IN THIS SECTION

- [Understanding Dual Configuration on P2-10G-40G-QSFPP PIC | 104](#)
- [Understanding Port Group | 105](#)
- [Port Numbering on P2-10G-40G-QSFPP PIC When Port Groups Are Not Configured | 109](#)
- [10-Gigabit Ethernet Mode | 112](#)
- [40-Gigabit Ethernet Mode | 113](#)

All the ports on the P2-10G-40G-QSFPP PIC are plugged into quad small form-factor pluggable plus transceivers (QSFP+) that, in turn, are connected to fiber-optic cables that support both 10-Gigabit Ethernet standards and 40-Gigabit Ethernet standards, thereby enabling you to configure the PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode.

You can configure the ports on the PIC in 10-Gigabit Ethernet mode or 40-Gigabit Ethernet mode at the port group level.

The following sections describe the P2-10G-40G-QSFPP PIC and the various framing modes that are supported on it:

Understanding Dual Configuration on P2-10G-40G-QSFPP PIC

All the ports on the P2-10G-40G-QSFPP PIC are QSFP+ based—that is, all the ports are connected to fiber-optic cables by means of QSFP+ transceivers.

The QSFP+ module—which includes the transceiver and the fiber-optic cable—supports the following standards on the P2-10G-40G-QSFPP PIC:

- 10-Gigabit Ethernet in LAN PHY framing mode (also known as native Ethernet mode) and WAN PHY framing mode.

Note that the ports follow a 4-level interface-naming convention—*et-fpc/pic/QSFP+ port:channel/in* this mode.

- 40-Gigabit Ethernet in LAN PHY framing mode.

Note that the ports follow a 3-level interface-naming convention—*et-fpc/pic/QSFP+ port* in this mode.

The P2-10G-40G-QSFPP PIC provides forty-eight 10-Gigabit Ethernet ports or twelve 40-Gigabit Ethernet ports. or .

The PIC can be configured either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode with the `set chassis fpc fpc-number pic pic-number pic-mode (10G | 40G)` configuration command. By default, the PIC is configured in 10-Gigabit Ethernet LAN PHY framing mode.

If you want to configure the PIC in 10-Gigabit Ethernet mode to operate in 40-Gigabit Ethernet mode, you must:

1. Delete all the interfaces in the PIC at the [edit interfaces] hierarchy level.
2. Configure the PIC to operate in 40-Gigabit Ethernet mode by using the `set chassis fpc fpc-slot pic pic-slot pic-mode 40G` configuration command and commit.

The PIC reboots and starts operating in the new mode.

The same procedure is applicable when you can configure the PIC in 40-Gigabit Ethernet PIC to operate in 10-Gigabit Ethernet mode. In this case, you must execute the `set chassis fpc fpc-slot pic pic-slot pic-mode 10G` configuration mode command.

To check the current diagnostics of the PIC, you must run the relevant operational mode CLI commands such as `show chassis hardware`, `show interfaces diagnostics optics interface-name`,

Understanding Port Group

The FPC2-PTX-P1A FPC can host two PICs and has eight Packet Forwarding Engines. The first four Packet Forwarding Engines on the FPC are associated with PIC 0 and the next four are associated with PIC 1.

All ports associated to one Packet Forwarding Engine compose a port group. Each PIC supports four Packet Forwarding Engines. Therefore, four port groups exist for each P2-10G-40G-QSFPP PIC.

Each Packet Forwarding Engine provides throughput of 120 Gbps.

Points to Remember

Consider the following points when configuring the PIC at the port group level:

- You can configure the ports in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode at the port group level.
- You can configure the port speed only on the first port in the port group. That is, you must configure the port speed for the port group on the ports numbered 0, 3, 6, and 9—the first ports in the respective port groups. An error message is logged when you try to configure the speed on any other port in the port group and this configuration will not have any effect on the PIC.
- You can configure the port speed of a port group only when the value of the `pic-mode` statement at the `[edit chassis fpc fpc-slot pic pic-slot pic-mode]` is set to `10G` or when the statement is not configured.
- You cannot configure different speeds for the ports in the same port group.
- You can configure different speeds for different port groups.

Port Group in 10-Gigabit Ethernet Mode

Each Packet Forwarding Engine supports twelve 10-Gigabit Ethernet ports in LAN PHY or in WAN PHY framing mode.

Note that when a port group is configured from 10-Gigabit Ethernet mode to 40-Gigabit Ethernet mode, the ports with 4-level interface-naming convention are deleted and three 40-Gigabit Ethernet mode ports with 3-level interface-naming convention are created.

Note that when the configuration of a port group is changed from 10-Gigabit Ethernet mode to 40-Gigabit Ethernet mode, the configuration of the twelve 10-Gigabit Ethernet ports is deleted and the 4-level interface-naming convention of the ports is also lost. Instead, three 40-Gigabit Ethernet ports are configured and these ports adhere to the 3-level interface-naming convention

Port Group in 40-Gigabit Ethernet Mode

Each Packet Forwarding Engine supports three 40-Gigabit Ethernet ports in LAN PHY framing mode.

Note that when the configuration of a port group is changed from 40-Gigabit Ethernet mode to 10-Gigabit Ethernet mode, the configuration of the three 40-Gigabit Ethernet ports is deleted and the 3-level interface-naming convention of the ports is also lost. Instead, twelve 10-Gigabit Ethernet ports are configured and these ports adhere to the 4-level interface-naming convention.

Port Number Mapping When Port Groups Are Configured

[Table 3 on page 106](#) shows the port numbering in 40-Gigabit Ethernet mode and in 10-Gigabit Ethernet mode at the port group level.

Table 3: Port Number Mapping When Port Groups Are Configured

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
0 (0)	et-1/1/0	et-1/1/0:0 et-1/1/0:1 et-1/1/0:2 et-1/1/0:3
	et-1/1/1	et-1/1/1:0 et-1/1/1:1 et-1/1/1:2 et-1/1/1:3
	et-1/1/2	et-1/1/2:0 et-1/1/2:1 et-1/1/2:2 et-1/1/2:3

Table 3: Port Number Mapping When Port Groups Are Configured *(Continued)*

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
3(1)	et-1/1/3	et-1/1/3:0 et-1/1/3:1 et-1/1/3:2 et-1/1/3:3
	et-1/1/4	et-1/1/4:0 et-1/1/4:1 et-1/1/4:2 et-1/1/4:3
	et-1/1/5	et-1/1/5:0 et-1/1/5:1 et-1/1/5:2 et-1/1/5:3
6(2)	et-1/1/6	et-1/1/6:0 et-1/1/6:1 et-1/1/6:2 et-1/1/6:3
	et-1/1/7	et-1/1/7:0 et-1/1/7:1 et-1/1/7:2 et-1/1/7:3

Table 3: Port Number Mapping When Port Groups Are Configured *(Continued)*

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
	et-1/1/8	et-1/1/8:0 et-1/1/8:1 et-1/1/8:2 et-1/1/8:3
6(2)	et-1/1/6	et-1/1/6:0 et-1/1/6:1 et-1/1/6:2 et-1/1/6:3
	et-1/1/7	et-1/1/7:0 et-1/1/7:1 et-1/1/7:2 et-1/1/7:3
	et-1/1/8	et-1/1/8:0 et-1/1/8:1 et-1/1/8:2 et-1/1/8:3
9(3)	et-1/1/9	et-1/1/9:0 et-1/1/9:1 et-1/1/9:2 et-1/1/9:3

Table 3: Port Number Mapping When Port Groups Are Configured (Continued)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
	et-1/1/10	et-1/1/10:0 et-1/1/10:1 et-1/1/10:2 et-1/1/10:3
	et-1/1/11	et-1/1/11:0 et-1/1/11:1 et-1/1/11:2 et-1/1/11:3

Port Numbering on P2-10G-40G-QSFPP PIC When Port Groups Are Not Configured

[Table 4 on page 109](#) shows the port numbering in 40-Gigabit Ethernet mode and in 10-Gigabit Ethernet mode when port groups are *not* configured on the P2-10G-40G-QSFPP PIC.

Table 4: Port Number Mapping When Port Groups Are Not Configured

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
0	et-1/1/0	et-1/1/0:0 et-1/1/0:1 et-1/1/0:2 et-1/1/0:3

Table 4: Port Number Mapping When Port Groups Are Not Configured (*Continued*)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
1	et-1/1/1	et-1/1/1:0 et-1/1/1:1 et-1/1/1:2 et-1/1/1:3
2	et-1/1/2	et-1/1/2:0 et-1/1/2:1 et-1/1/2:2 et-1/1/2:3
3	et-1/1/3	et-1/1/3:0 et-1/1/3:1 et-1/1/3:2 et-1/1/3:3
4	et-1/1/4	et-1/1/4:0 et-1/1/4:1 et-1/1/4:2 et-1/1/4:3
5	et-1/1/5	et-1/1/5:0 et-1/1/5:1 et-1/1/5:2 et-1/1/5:3

Table 4: Port Number Mapping When Port Groups Are Not Configured (*Continued*)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
6	et-1/1/6	et-1/1/6:0 et-1/1/6:1 et-1/1/6:2 et-1/1/6:3
7	et-1/1/7	et-1/1/7:0 et-1/1/7:1 et-1/1/7:2 et-1/1/7:3
8	et-1/1/8	et-1/1/8:0 et-1/1/8:1 et-1/1/8:2 et-1/1/8:3
9	et-1/1/9	et-1/1/9:0 et-1/1/9:1 et-1/1/9:2 et-1/1/9:3
10	et-1/1/10	et-1/1/10:0 et-1/1/10:1 et-1/1/10:2 et-1/1/10:3

Table 4: Port Number Mapping When Port Groups Are Not Configured (Continued)

QSFP+ Port Number	Port Numbering in 40-Gigabit Ethernet Mode	Port Numbering in 10-Gigabit Ethernet Mode
11	et-1/1/11	et-1/1/11:0 et-1/1/11:1 et-1/1/11:2 et-1/1/11:3

10-Gigabit Ethernet Mode

A 10-Gigabit Ethernet interface can operate in 10-Gigabit Ethernet LAN PHY framing mode or in 10-Gigabit Ethernet WAN PHY framing mode.

You can configure a 10-Gigabit Ethernet interface at the [edit interface *interface-name* framing-mode (lan-phy | wan-phy)] hierarchy level to operate in 10-Gigabit Ethernet LAN PHY framing mode or in 10-Gigabit Ethernet WAN PHY framing mode.

Each P2-10G-40G-QSFPP PIC provides 48 physical interfaces. The interfaces are represented by the 4-level interface-naming convention—et-*fpc/pic/QSFP+ port:channel*, where the value of the *QSFP+ port* option ranges from 0 through 11 and the value of the *channel* option ranges from 0 through 3.

Framing Mode Overview

When a P2-10G-40G-QSFPP PIC is configured in 10-Gigabit Ethernet framing mode, it can operate in one of the following framing modes:

- LAN PHY framing mode. Note that by default, the PIC is in 10-Gigabit Ethernet LAN PHY framing mode. You can configure loopback at the [edit interfaces *interface-name* sonet-options loopback] hierarchy level.

The ports are set to LAN PHY framing mode by default when the *framing-mode* statement is not configured at the [edit interface *interface-name*] hierarchy level.

- WAN PHY framing mode

Supported Features on LAN PHY and WAN PHY Framing Mode

The following features are supported in LAN PHY and WAN PHY framing mode when the PIC operates in 10-Gigabit Ethernet mode:

- The following are supported for WAN interface sublayer statistics, defects, and alarms when the PIC operates in WAN PHY framing mode:
 - GR 253 standard.
 - show interfaces *interfaces-name* operational mode command displays WAN interface sublayer statistics, defects and alarms.
 - Interrupt-driven notification for WAN interface sublayer defects.
 - Path trace and trigger options for WAN interface sublayer alarms.
 - Transmitting and receiving J1 (path trace) messages—J1 is a part of path overhead in a WAN interface sublayer frame.
 - Line loopback and local loopback. Loopback is configured at the [edit interfaces *interface-name* sonet-options loopback] hierarchy level in WAN PHY framing mode.
 - The defects PHY LOL (loss of light) and PHY PLL (loss of PLL lock) are detected and reported at the physical level in WAN PHY framing mode.

Fast reroute (FRR) in WAN PHY framing mode:

- Enable or disable preemptive fast reroute (FRR) options at the [edit interfaces *interface-name* otn-options preemptive-fast-reroute] hierarchy level.
- Configure thresholds and interval for the optical channel data unit (ODU) signal degradation (*odu-signal-degrade*) and the configurable pre-FEC bit error rate (BER) (*ber-threshold-signal-degrade*) at the [edit interfaces *interface-name* otn-options odu-signal-degrade] hierarchy level and the [edit interfaces *interface-name* otn-options signal-degrade] hierarchy level, respectively.

40-Gigabit Ethernet Mode

You can configure twelve 40-Gigabit Ethernet interfaces that operate in LAN PHY framing mode. The interfaces are represented by the 3-level interface-naming convention *et-fpc/pic/QSFP+ port*, where the value of the *QSFP+ port* variable ranges from 0 through 11.

Configure the P2-10G-40G-QSFPP PIC

IN THIS SECTION

- [Configure the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 114](#)
- [Configure the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 115](#)
- [Configure the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 115](#)
- [Configure the PIC at Port Group Level | 116](#)
- [Configure Framing Mode on P2-10G-40G-QSFPP PIC | 117](#)

Starting with Junos OS Release 14.1R2, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC. You can configure the P2-10G-40G-QSFPP PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode..

The following tasks explain how to configure the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode and to configure the framing modes on it.

Configure the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode:

1. In configuration mode, go to the [edit chassis] hierarchy level.

```
[edit]
user@host# edit chassis
```

2. Configure the PIC in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode after specifying the required FPC slot and PIC slot. Note that all the PIC ports in a PIC are configured at once with this configuration command.

```
[edit chassis]
user@host# set fpc fpc-slot pic pic-slot pic-mode (10G | 40G)
```

Configure the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC that is configured in 10-Gigabit Ethernet mode to operate in 40-Gigabit Ethernet mode:

1. In configuration mode, go to the [edit interfaces] hierarchy level.

```
[edit]
user@host# edit interfaces
```

2. Delete all the interfaces in the PIC, commit, and then move to the top of the hierarchy level.

```
[edit interfaces]
user@host# delete interface-name
user@host# delete ...
user@host# commit
user@host# top
```

3. Configure the PIC to operate in 40-Gigabit Ethernet mode and commit.

```
[edit]
user@host# set chassis fpc fpc-slot pic pic-slot pic-mode 40G
user@host# commit
```

After the configuration is committed, the PIC reboots and starts operating in the 40-Gigabit Ethernet mode. You can now configure the parameters, such as encapsulation, framing mode, and so on, for the twelve 40-Gigabit Ethernet interfaces in the PIC as needed.

Configure the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC that is configured in 40-Gigabit Ethernet mode to operate in 10-Gigabit Ethernet mode:

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

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

2. Delete all the interfaces in the PIC, commit, and then move to the top of the hierarchy level.

```
[edit interfaces]
user@host# delete interface-name
user@host# delete ...
user@host# commit
user@host# top
```

3. Configure the PIC to operate in 10-Gigabit Ethernet mode and commit.

```
[edit]
user@host# set chassis fpc fpc-slot pic pic-slot pic-mode 10G
user@host# commit
```

After the configuration is committed, the PIC reboots and starts operating in the 10-Gigabit Ethernet mode. You can now configure the parameters, such as encapsulation, framing mode, and so on, for the forty-eight 10-Gigabit Ethernet interfaces in the PIC as needed.

Configure the PIC at Port Group Level

Before You Begin

Verify that the `pic-mode` statement at the `[edit chassis fpc fpc-slot pic pic-slot pic-mode]` is not configured or that its value is set to `10G`.

To configure a port group in the P2-10G-40G-QSFPP PIC to operate in 10-Gigabit Ethernet mode or 40-Gigabit Ethernet mode:

1. In configuration mode, go to the `[edit chassis fpc fpc-slot pic pic-slot]` hierarchy level.

```
[edit]
user@host# edit chassis fpc fpc-slot pic pic-slot
```

2. Configure the port number as 0, 3, 6, or 9 and the speed as `10G` or `40G`. Note that you can configure the port speed only on the first port in the port group. That is, configure the port speed only on the ports numbered 0, 3, 6, and 9. An error message is displayed when you try to configure the speed on any other port in the port group.

```
[edit chassis fpc fpc-slot pic pic-slot]
user@host# set port port-number speed (10G | 40G)
```

A system log message is logged when you try to configure a different port speed on a port when the port group is operating at another speed.

Configure Framing Mode on P2-10G-40G-QSFPP PIC

IN THIS SECTION

- [Configure LAN PHY or WAN PHY Framing Mode in 10-Gigabit Ethernet Mode | 117](#)
- [Configure LAN PHY Framing Mode in 40-Gigabit Ethernet Mode | 118](#)

You can configure LAN PHY, or WAN PHY framing mode when the PIC is operating in 10-Gigabit Ethernet mode. You can configure LAN PHY framing mode when the PIC is operating in 40-Gigabit Etherent mode.

Configure LAN PHY or WAN PHY Framing Mode in 10-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode to operate in LAN PHY framing mode or in WAN PHY framing mode, you must configure the framing mode individually on all the interfaces:

You can configure LAN PHY, or WAN PHY framing mode when the PIC is operating in 10-Gigabit Ethernet mode. You can configure LAN PHY framing mode when the PIC is operating in 40-Gigabit Etherent mode. The following tasks explain how to configure the various framing modes on the PIC

1. In configuration mode, go to the [edit interfaces *interfaces-name*] hierarchy level, where the interface name is in *et-fpc/pic/port:channel* format.

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

2. Configure the framing mode as LAN PHY or WAN PHY and commit.

```
[edit interfaces interface-name]
user@host# set framing (lan-phy | wan-phy)
user@host# commit
```

For example, you can configure the framing mode as LAN PHY or WAN PHY on the *et-1/1/1:0* interface.

Configure LAN PHY Framing Mode in 40-Gigabit Ethernet Mode

To configure the P2-10G-40G-QSFPP PIC in 40-Gigabit Ethernet mode to operate in LAN PHY framing mode:

1. In configuration mode, go to the [edit interfaces *interfaces-name*] hierarchy level, where the interface name is in et-*fpc/pic/port* format.

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

2. Configure the framing mode as LAN PHY and commit.

```
[edit interfaces interface-name]
user@host# set framing (lan-phy)
user@host# commit
```

For example, you can configure the framing mode as LAN PHY on the et-2/2/2 interface.

Example: Configure the P2-10G-40G-QSFPP PIC

IN THIS SECTION

- Requirements | [118](#)
- Overview | [119](#)
- Configuration | [119](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 14.1R2 or Junos OS Release 14.2 or later
- One PTX5000 router with P2-10G-40G-QSFPP PIC

Overview

Starting with Junos OS Release 14.1R2 and 14.2R1, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC.

All the ports on the P2-10G-40G-QSFPP PIC are QSFP+ based—that is, all the ports are connected to fiber-optic cables by means of QSFP+ transceivers. The P2-10G-40G-QSFPP PIC provides forty-eight 10-Gigabit Ethernet ports or twelve 40-Gigabit Ethernet ports.

The QSFP+ module—which includes the transceiver and the fiber-optic cable—supports the following standards on the P2-10G-40G-QSFPP PIC:

- 10-Gigabit Ethernet in LAN PHY framing mode (also known as native Ethernet mode) and WAN PHY framing mode.
- 40-Gigabit Ethernet in LAN PHY framing mode.

Configuration

IN THIS SECTION

- [Verification | 121](#)

To configure the P2-10G-40G-QSFPP PIC to operate in 10-Gigabit Ethernet mode, and to set the framing mode and other options on an interface on this PIC, perform the following tasks:

Configuring the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet Mode

Step-by-Step Procedure

Configure the PIC in 10-Gigabit Ethernet mode.

1. In configuration mode, go to the [edit chassis] hierarchy level.

```
[edit]
user@host# edit chassis
```

- Configure the PIC in 10-Gigabit Ethernet mode after specifying the required FPC slot and PIC slot. Note that the PIC restarts after the configuration is committed and all the ports in the PIC come up in the 10-Gigabit Ethernet mode.

```
[edit chassis]
user@host# set fpc 1 pic 1 pic-mode 10G
```

Configuring the Framing Mode on an Interface

Step-by-Step Procedure

To configure an interface *et-1/1/1:0* in the P2-10G-40G-QSFPP PIC to operate in LAN PHY framing mode:

- In configuration mode, go to the [edit interfaces et-1/1/1:0] hierarchy level.

```
[edit]
user@host# edit interfaces et-1/1/1:0
```

- Configure the framing mode for the interface as LAN PHY and commit.

```
[edit interfaces et-1/1/1:0]
user@host# set framing lan-phy
user@host# commit
```

Similarly, you can configure LAN PHY or WAN PHY framing mode for the other interfaces in the PIC.

Configuring the Interface Options

Step-by-Step Procedure

Configure the interface options for the interface et-1/1/1:0 as needed. The following procedure configures a few interface-specific options.

- In configuration mode, go to the [edit interfaces et-1/1/1:0] hierarchy level.

```
[edit]
user@host# edit interfaces et-1/1/1:0
```

2. Configure the encapsulation as ethernet-ccc.

```
[edit interfaces et-1/1/1:0]
user@host# set encapsulation ethernet-ccc
```

3. Configure the family as CCC for the logical interface 0.

```
[edit interfaces et-1/1/1:0]
user@host# set unit 0 family ccc
```

4. Enable flow control to regulate the flow of packets from the router to the remote side of the network connection.

```
[edit interfaces et-1/1/1:0 gigether-options]
user@host# set flow-control
```

5. Enable loopback mode for the interface, commit the configuration, and exit the configuration mode.

```
[edit interfaces et-1/1/1:0 gigether-options]
user@host# set loopback
user@host# commit
user@host# quit
```

Verification

IN THIS SECTION

- [Displaying Interface Details | 121](#)

Displaying Interface Details

Purpose

To display interface-specific details of the et-1/1/1:0 interface.

Action

Execute the `show interfaces et-1/1/1:0` operational command.

```
user@host# run show interfaces et-1/1/1:0
Interface index: 525, SNMP ifIndex: 522
  Link-level type: Ethernet, MTU: 1514, MRU: 0, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags : None
  CoS queues : 8 supported, 8 maximum usable queues
  Current address: ac:4b:c8:f6:af:68, Hardware address: ac:4b:c8:f6:af:68
  Last flapped : 2014-07-25 02:23:56 PDT (02:16:07 ago)
  Input rate : 0 bps (0 pps)
  Output rate : 0 bps (0 pps)
  Active alarms : LINK
  Active defects : LINK
  PCS statistics          Seconds
    Bit errors             0
    Errorred blocks        1
  Interface transmit statistics: Disabled
```

Meaning

The interface details are displayed. Note that to display information for an interface in 10-Gigabit Ethernet mode for the P2-10G-40G-QSFPP PIC, you must use the `et-fpc/pic/port.channel/format`.

Framing Overview

The 10-Gigabit Ethernet interfaces support operation in two modes:

- 10GBASE-R, LAN Physical Layer Device (LAN PHY)
- 10GBASE-W, WAN Physical Layer Device (WAN PHY)

When the external interface is running in LAN PHY mode, it bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface. When the external

interface is running in WAN PHY mode, it uses the WIS sublayer to transport 10-Gigabit Ethernet frames in an OC192c SONET payload.

WAN PHY mode is supported on MX240, MX480, MX960, and PTX Series Packet Transport routers only.

Although the external interface provides a lower throughput when running in WAN PHY mode because of the extra SONET overhead, it can interoperate with SONET section or line level repeaters. This creates an advantage when the interface is used for long-distance, point-to-point 10-Gigabit Ethernet links. When the external interface is running in WAN PHY mode, some SONET options are supported.

Understanding WAN Framing

If you use the `wan-phy` statement option at the `[edit interfaces xe-fpc/pic/0 framing]` hierarchy level to configure Trio WAN mode framing for 10-Gigabit Ethernet interfaces, then the alarm behavior of the link, although in full compliance with the IEEE 802.3ae 10-Gigabit Ethernet standard, might not be as expected.

In particular:

- The interface does not distinguish between loss of light (LOL), loss of phase lock loop (PLL), or loss of signal (LOS). If a loss of PLL or LOS alarm occurs, then both PLL and LOS alarms are raised. LOL is also raised because there is no separate LOL indication from the hardware.
- The interface does not raise LOS, PLL, or LOL alarms when the fiber is disconnected from the interface port. You must remove the hardware to raise this alarm.
- The interface line-level alarm indicator signal (AIS-L) is not always raised in response to a loss of framing (LOF) defect alarm.
- If the AIS-L or path-level AIS (AIS-P) occurs, the interface path-level loss of code delineation (LCD-P) is not detected. LCD-P is seen during the path-level remote defect indicator (RDI-P) alarm.
- If an AIS-L alarm occurs, the AIS-P is not detected, but the LOP alarm is detected.

None of the alarm issues are misleading, but they make troubleshooting the root cause of problems more complex.

Configure Ethernet Framing

The 10-Gigabit Ethernet interfaces uses the interface type `xe-fpc/pic/port`. On single port devices, the port number is always zero.

The `xe-fpc/pic/port` interface inherits all the configuration commands that are used for gigabit Ethernet (`ge-fpc/pic/port`) interfaces.

To configure LAN PHY or WAN PHY operating mode, include the `framing` statement with the `lan-phy` or `wan-phy` option at the `[edit interfaces xe-fpc/pic/0]` hierarchy level.

```
[edit interfaces xe-fpc/pic/0 framing]
framing (lan-phy | wan-phy);
```

- On PTX Series Transport Routers, WAN PHY mode is supported only on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+
- When the PHY mode changes, interface traffic is disrupted because of port reinitialization.

To display interface information, use the operational mode command `show interfaces xe-fpc/pic/port extensive`.

Gigabit Ethernet MICs

SUMMARY

Learn about the Gigabit Ethernet MICs in this topic. You can configure interoperability between two 100-Gigabit Ethernet PICs.

IN THIS SECTION

- [100-Gigabit Ethernet Interfaces Interoperability | 125](#)
- [Platform-Specific 100-Gigabit Ethernet Interfaces Behavior | 125](#)

A 100 Gigabit Ethernet Interface is a high-speed networking technology that allows data transmission at a rate of 100 Gbps, providing significantly faster connectivity compare with traditional Ethernet standards

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Review the ["Platform-Specific 100-Gigabit Ethernet Interfaces Behavior" on page 125](#) section for notes related to your platform

100-Gigabit Ethernet Interfaces Interoperability

Juniper Networks Junos operating system (Junos OS) supports a variety of 100-Gigabit Ethernet interfaces. The 100-Gigabit Ethernet standard, introduced by IEEE 802.3ba-2010, enables transmission of Ethernet frames at the rate of 100 gigabits per second (Gbps). It is used for very high speed transmission of voice and data signals across the numerous world-wide fiber-optic networks.

Interface interoperability refers to the ability of an interface to interoperate with other router interfaces. You can enable interoperability between different 100-Gigabit Ethernet interfaces by performing specific configuration tasks. The following sections list the 100-Gigabit Ethernet interfaces, corresponding interoperable interfaces, and links to the interoperability tasks and reference information.

SEE ALSO

[Periodic Packet Management on Ethernet Interfaces | 93](#)

Platform-Specific 100-Gigabit Ethernet Interfaces Behavior

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platform:

Platform	Difference
MX Series	<ul style="list-style-type: none"> MX240, MX480, MX960, MX2010, and MX2020 routers that support 100-Gigabit Ethernet Interfaces, support the following interface modules: <ul style="list-style-type: none"> 100-Gigabit Ethernet MIC with CFP 100-Gigabit Ethernet MIC with CXP 100-Gigabit Ethernet ports on the MPC4E MX2010 and MX2020 routers that support 100-Gigabit Ethernet Interfaces, support the following interface modules: <ul style="list-style-type: none"> 100-Gigabit Ethernet MIC with CFP2 100-Gigabit Ethernet MIC with CXP (4 Ports)
PTX Series	<ul style="list-style-type: none"> PTX5000 routers that support 100-Gigabit Ethernet Interfaces, support the following PICs: <ul style="list-style-type: none"> 100-Gigabit Ethernet PIC with CFP 100-Gigabit Ethernet PIC with CFP2 100-Gigabit Ethernet OTN PIC PTX 3000 and PTX5000 routers that support 100-Gigabit Ethernet Interfaces, support the following PIC: <ul style="list-style-type: none"> 100-Gigabit DWDM OTN PIC

Smart SFPs To Transport Legacy Network Traffic over Packet Switched Networks

SUMMARY

Learn how to transport legacy TDM traffic over Packet switched networks using Smart SFP transceivers.

IN THIS SECTION

- [Smart SFPs To Transport Legacy Traffic | 127](#)
- [Smart-SFPs for transporting legacy PDH Traffic | 130](#)
- [Configure the Smart-SFPs to transport legacy SDH Traffic | 136](#)

Smart SFPs To Transport Legacy Traffic

IN THIS SECTION

- [Smart SFP Transceivers for Transporting PDH Traffic over PSNs Overview | 128](#)
- [Smart SFP Transceivers for Transporting SDH Traffic over PSNs Overview | 128](#)
- [Benefits of Smart SFP Transceivers | 129](#)

Legacy networks such as SONET and SDH, which are used for very high-speed transmission of voice and data signals across the numerous fiber-optic networks, still operate worldwide. These legacy networks use time-division multiplexing (TDM), which ensures that a constant stream of data travels on the network. Lower bit-rate streams of information are combined, or multiplexed, up into higher bit-rate streams to take advantage of the bandwidth available. Today, as data is the most significant type of traffic on the legacy networks, most organizations are planning to migrate their existing legacy networks to packet-switched networks (PSNs), which are better suited for data transport. However, a part of the network traffic continues to remain TDM-based. And migrations are expensive and require detailed planning for allocation of rack space, power, and new equipment.

To ensure seamless migration from legacy networks to PSNs in a cost-effective and space-optimized manner, you can use smart small form-factor pluggable (SFP) transceivers. Install a smart SFP transceiver on your router or switch and easily transport TDM traffic (converted into a packet stream) across a PSN.

TDM traffic is broadly classified into: *Plesiochronous Digital Hierarchy* (PDH) and *Synchronous Digital Hierarchy* (SDH) traffic. Both PDH and SDH technologies are associated with digital multiplexers. In PDH traffic, the bit streams are of same bit rate but are derived from different clocks that belong to different oscillators. Hence, the name *Plesiochronous*.

Smart SFP Transceivers for Transporting PDH Traffic over PSNs Overview

Junos OS supports the following three smart SFP transceivers on routers for transporting PDH traffic over PSNs:

- DS3 smart SFP (SFP-GE-TDM-DS3)
- E1 smart SFP (SFP-GE-TDM-E1)
- T1 smart SFP (SFP-GE-TDM-T1)

The Gigabit Ethernet MIC with SFP (MIC-3D-20GE-SFP), Gigabit Ethernet MIC with SFP (E) (MIC-3D-20GE-SFP-E), Gigabit Ethernet MIC with SFP (EH) (MIC-3D-20GE-SFP-EH), and the Gigabit Ethernet MIC with 256b-AES MACSEC (MIC-MACSEC-20GE) modular interface cards support the smart SFP transceivers.

The smart SFP transceivers encapsulate traffic on PDH interfaces on the WAN side as Ethernet frames on the system side. The encapsulated traffic from the transceivers is sent over the PSNs provisioned across the network. You can further encapsulate the PDH interfaces using MEF8 or MPLS framing. You can also configure single or dual VLAN tagging within a default emulation circuit identifier (ECID).

You can encapsulate E1 and T1 traffic as Ethernet frames by using Structure Agnostic TDM over Packet (SAToP) using MEF8 framing. MPLS framing and both single and dual VLAN tagging are supported. E1 and T1 encapsulation uses the SAToP according to Transparent PDH over Packet (TPoP) standard. You can encapsulate DS3 traffic as Ethernet frames using MEF8 or MPLS framing. Both single and dual VLAN tagging is supported. DS3 encapsulation uses the Virtual Container over Packet (VCoP) standard.

At the local end, the smart SFP transceiver slices the TDM data stream, encapsulates the Ethernet frames and pushes it onto the PSN. The smart SFPs are always paired on the other end of the emulated circuit, and are preconfigured to be in the same multicast MAC address group. At the far end, the smart SFP transceiver decapsulates the Ethernet frames, re-builds the TDM data stream, and forwards it onto the local TDM interface.

Smart SFP Transceivers for Transporting SDH Traffic over PSNs Overview

Junos OS supports the following three smart SFP transceivers on MX Series routers for transporting SDH traffic over PSNs:

- STM1 smart SFP (SFP-GE-TDM-STM1)
- STM4 smart SFP (SFP-GE-TDM-STM4)

- STM16 smart SFP (SFP-GE-TDM-STM16)

On your router, the MPC1, MPC2, MPC3 line cards, the Gigabit Ethernet MIC with SFP (MIC-3D-20GE-SFP), Gigabit Ethernet MIC with SFP (E) (MIC-3D-20GE-SFP-E), Gigabit Ethernet MIC with SFP (EH) (MIC-3D-20GE-SFP-EH), and the Gigabit Ethernet MIC with 256b-AES MACSEC (MIC-MACSEC-20GE) modular interface cards support the smart SFP transceivers. Only the 10-Gigabit Ethernet interfaces on the 256-AES MACSEC MIC support the STM16 smart SFP (SFP-GE-TDM-STM16). The MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE) line card supports the STM16 smart SFP transceiver.

The smart SFP transceivers encapsulate traffic on SDH interfaces on the WAN side as Ethernet frames on the system side. The encapsulated traffic from the SFP transceivers is sent over the PSNs provisioned across the network. You can encapsulate the SDH interfaces using MEF8 framing. You can also configure single VLAN tagging within a default emulation circuit identifier (ECID).

You can encapsulate STM traffic as Ethernet frames using MEF8 framing. Only single VLAN tagging is supported. STM encapsulation uses the Transparent SONET/SDH over Packet (TSoP) standard.

At the local end, the smart SFP transceiver slices the TDM data stream, encapsulates the Ethernet frames and pushes it onto the PSN. The smart SFPs are always paired on the other end of the emulated circuit, and are preconfigured to be in the same multicast MAC address group. At the far end, the smart SFP transceiver decapsulates the Ethernet frames, re-builds the TDM data stream and forwards it onto the local TDM interface.

Benefits of Smart SFP Transceivers

- Lower operational costs—Smart SFP transceivers enable easy and simplified migration and upgrades from legacy networks to PSNs.
- Operational simplicity and flexibility—You don't need to configure individual TDM interfaces over packet connections. You can deploy additional equipment only if you need it.
- Space saving. Does not require additional rack space.
- Low carbon footprint. Lower power consumption and existing equipment such as TDM access nodes are still in use after migration. Reduced electronic waste.
- Definitive migration to only-Ethernet based network equipment, removing the need to have dedicated TDM network interface cards for the termination of the TDM lines.

Smart-SFPs for transporting legacy PDH Traffic

IN THIS SECTION

- [Requirements | 130](#)
- [Overview | 130](#)
- [Configure the DS3 Smart SFP | 130](#)
- [Verification | 133](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 19.4R1 or later
- A single router
- DS3 Smart SFP (SFP-GE-TDM-DS3) transceiver

Overview

This example provides information about configuring the DS3 Smart SFP (SFP-GE-TDM-DS3) transceiver on a router to enable the transceiver to encapsulate DS3 packets as Ethernet frames while transporting the packets from legacy networks to PSNs. You can configure the DS3 Smart SFP to further encapsulate the DS3 packets using MEF8 or MPLS framing as VCoP (Virtual container over Packet) for DS3 interfaces. You can also configure single or dual VLAN tagging.

The TDM data stream, sliced, and encapsulated into Ethernet frames is pushed into the PSN to reach the far end point of a similar SFP type. The Smart SFPs are always paired on the other end of the emulated circuit and are pre-configured to be part of the same multicast MAC address group. At the far end, the smart SFP transceiver decapsulates the Ethernet frames, re-builds the TDM data stream and forwards it onto the local TDM interface.

Configure the DS3 Smart SFP

IN THIS SECTION

- [Procedure | 131](#)

Procedure

Step-by-Step Procedure

In this example, you configure the Smart SFP to transport PDH traffic over PSN networks. To configure the Smart SFP, perform the following tasks:

1. In Configuration mode, create a valid Interface to enable the Smart SFP to communicate with the Junos OS. Configuring VLAN tagging creates a control interface.

```
[edit]
user@host #set interfaces ge-4/0/0 unit 0
user@host #set interfaces ge-4/0/0 vlan-tagging
```

2. Specify the type of Smart SFP to be configured on the interface. In this example, we are configuring a DS3 Smart SFP.

```
[edit]
user@host #set interfaces ge-4/0/0 tdm-options sfp-type DS3
```

3. (Optional) Configure the destination MAC address on the local end smart SFP using the `dmac-address` statement at the [edit interfaces ge-4/0/0 tdm-options] hierarchy level to encapsulate the MAC address of the far end smart SFP. To enable MAC address validation or checking of the destination MAC address at the far end smart SFP, use the `ces-psn-port-dmac-check-enable` statement. If the MAC address of the packet does not match, the packet is discarded.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-channel dmac-address dmac-address
```

4. (Optional) Configure the encapsulation mode (MEF8 or MPLS) for further network processing. The default encapsulation mode for DS3 Smart SFP is MEF8.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-channel mode mode
```

5. (Optional) Configure single or dual VLAN tagging on the encapsulated packets. DS3 Smart SFP supports both single and dual VLAN tagging. If you want to configure single VLAN tagging, use the `vlan-id-1` statement and specify the VLAN ID. If you want to configure dual VLAN tagging, use `vlan-`

`id-1` and `vlan-id-2` statements to configure the inner and outer VLAN IDs. Possible values for the VLAN ID: 0 through 4094.

Single VLAN tagging

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-channel vlan-id-1 vlan-id
```

Dual VLAN tagging

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-channel vlan-id-1 vlan-id
user@host # set ces-psn-channel vlan-id-2 vlan-id
```

6. (Optional) Configure the emulation circuit ID for encapsulation and decapsulation. If you do not specify an emulation circuit ID, the default value is 0. Possible values for the encapsulation and decapsulation ID: 0 through 1048575.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set iwf-params encap-ecid encap-ecid
user@host # set iwf-params decap-ecid decap-ecid
```

7. (Optional) Specify if you require checking of the destination MAC address of the incoming packets on the receiving SFP at the [edit interfaces ge-4/0/0 tdm-options] hierarchy. If you have configured the destination MAC address using the `dmac-address` option, use this option to verify the MAC address on the receiving SFP. If you have enabled MAC address verification and the MAC address does not match, the packet is discarded by the smart SFP.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-port-dmac-check-enable
```

8. (Optional) Enable looping back of the input path of TDM traffic on the SFP TDM port. The input path refers to the traffic from the TDM side that is looped back.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set tdm-in-loop
```

9. (Optional) Enable looping back of the output path of TDM traffic on the SFP TDM port. The output path refers to the traffic from the Ethernet side that is looped back.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set tdm-out-loop
```

Verification

IN THIS SECTION

- [Verifying the DS3 Smart SFP Statistics on the Interface | 133](#)
- [Verifying the DS3 Smart SFP Defects on the Interface | 135](#)

To verify that the DS3 Smart SFP is configured on the router, perform the following tasks:

Verifying the DS3 Smart SFP Statistics on the Interface

Purpose

To verify that the DS3 Smart SFP is configured on the router and to view the DS3 Smart SFP statistics.

Action

To view the DS3 Smart SFP statistics on the Interface, use the `show interfaces ge-4/0/0 smart-sfp-statistics` command.

```
user@host > show interfaces ge-4/0/0 smart-sfp-statistics
Physical interface: ge-4/0/0, Enabled, Physical link is Up
  Interface index: 281, SNMP ifIndex: 742
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
  None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags  : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
```

```

Link flags      : None
Smart Transceiver Type: DS3
Smart SFP: Configurable SFP
Smart SFP Ethernet port[P1] Statistics:          Counters
  Rx frames          1187126
  Rx bytes           149855236
  Rx errored fcs frames      0
  Rx unicast frames      1187124
  Rx multicast frames      2
  Rx broadcast frames      0
  Rx fragments          0
  Rx undersize frames      0
  Rx oversize frames      0
  Rx invalid vlan mismatch frames      0
  Tx frames           1392780998
  Tx bytes            1796396824
  Tx unicast frames      1377974
  Tx multicast frames      1391403024
  Tx broadcast frames      0
Smart SFP AV IWF Encap/Decap Statistics:          Counters
  Rx Packets          0
  Tx Packets          0
  Malformed Packets      0
  Reordered Packets      0
  Misordered Dropped Packets      0
  Missing Packets      0
  PlayedOut Packets      0
  JitterBuffer Overrun      0
  JitterBuffer Underrun      0
Smart SFP DS3 port[P0] statistics:          Counters
  BiPolarVariations/Excessive zero errors      0
  Tx B3 Errors          0
  Code Violation path errors      0
Logical interface ge-4/1/0.0 (Index 350) (SNMP ifIndex 615)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Input packets : 1384454023
  Output packets: 0
  Protocol multiservice, MTU: Unlimited

```

Meaning

The DS3 Smart SFP is configured on the router and you can view the DS3 Smart SFP statistics.

Verifying the DS3 Smart SFP Defects on the Interface

Purpose

To verify that the DS3 Smart SFP is configured on the router and to view the DS3 Smart SFP defects.

Action

To view the DS3 Smart SFP defects on the Interface, use the `show interfaces ge-4/0/0 smart-sfp-defects` command.

```
user@host > show interfaces ge-4/0/0 smart-sfp-defects
Physical interface: ge-4/0/0, Enabled, Physical link is Up
  Interface index: 211, SNMP ifIndex: 735
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error: None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags : None
  Smart Transceiver Type: DS3
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    Encapsulation Circuit Id: 16
    Decapsulation Circuit Id: 16
    Mode: MPLS
  Smart SFP Defects:
    TDM defects : Loss of Signal, Local Packet Loss, Out Of Frame
    System defects : None

Logical interface ge-4/0/0.0 (Index 355) (SNMP ifIndex 612)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Input packets : 430
  Output packets: 0
  Protocol multiservice, MTU: Unlimited
```

Meaning

The DS3 Smart SFP is configured on the router and you can view the DS3 Smart SFP defects.

Configure the Smart-SFPs to transport legacy SDH Traffic

IN THIS SECTION

- [Requirements | 136](#)
- [Overview | 136](#)
- [Configure the STM1 Smart SFP | 137](#)
- [Verification | 139](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 19.4R1 or later
- A single router
- STM1 Smart SFP (SFP-GE-TDM-STM1) transceiver

Overview

This example provides information about configuring the STM1 Smart SFP (SFP-GE-TDM-STM1) transceiver on a router to enable the transceiver to encapsulate STM1 packets as Ethernet frames while transporting the packets from legacy networks to PSNs. You can configure the STM1 Smart SFP to further encapsulate the STM1 packets using MEF8 framing as TSoP (Transparent SONET/SDH over Packet) for STM1 interfaces. You can only configure single VLAN tagging.

The TDM data stream, sliced, and encapsulated into Ethernet frames is pushed into the PSN to reach the far end point of a similar SFP type. The Smart SFPs are always paired on the other end of the emulated circuit and are pre-configured to be part of the same multicast MAC address group. At the far end, the smart SFP transceiver decapsulates the Ethernet frames, re-builds the TDM data stream and forwards it onto the local TDM interface.

Configure the STM1 Smart SFP

IN THIS SECTION

- [Procedure | 137](#)

Procedure

Step-by-Step Procedure

In this example, you configure the Smart SFP to transport SDH packets over PSNs. To configure the Smart SFP, perform the following tasks:

1. In Configuration mode, create a valid Interface to enable the Smart SFP to communicate with the Junos OS. Configuring VLAN tagging creates a control interface.

```
[edit]
user@host #set interfaces ge-3/0/0 unit 0
user@host #set interfaces ge-3/0/0 vlan-tagging
```

2. Specify the type of Smart SFP to be configured on the interface. In this example, we are configuring a STM1 Smart SFP.

```
[edit]
user@host #set interfaces ge-3/0/0 tdm-options sfp-type STM1
```

3. (Optional) Configure the destination MAC address using the `dmac-address` statement at the [edit interfaces ge-3/0/0 tdm-options] hierarchy level to encapsulate the MAC address of the far end smart SFP. To enable MAC address validation or checking of the destination MAC address at the far end smart SFP, use the `ces-psn-port-dmac-check-enable` statement. If the MAC address of the packet does not match, the packet is discarded.

```
[edit interfaces ge-3/0/0 tdm-options]
user@host # set ces-psn-channel dmac-address dmac-address
```

4. (Optional) Configure the encapsulation mode (MEF8 only) for further network processing. The default encapsulation mode for STM1 Smart SFP is MEF8.

```
[edit interfaces ge-3/0/0 tdm-options]
user@host # set ces-psn-channel mode mode
```

5. (Optional) Configure single VLAN tagging on the encapsulated packets. STM1 Smart SFP supports only single VLAN tagging. If you want to configure single VLAN tagging, use the *vlan-id-1* statement and specify the VLAN ID. Possible values for the VLAN ID: 0 through 4094.

Single VLAN tagging

```
[edit interfaces ge-3/0/0 tdm-options]
user@host # set ces-psn-channel vlan-id-1 vlan-id
```

6. (Optional) Configure the emulation circuit ID for encapsulation and decapsulation. If you do not specify an emulation circuit ID, the default value is 0. Possible values for the encapsulation and decapsulation ID: 0 through 1048575.

```
[edit interfaces ge-3/0/0 tdm-options]
user@host # set iwf-params encap-ecid encap-ecid
user@host # set iwf-params decap-ecid decap-ecid
```

7. (Optional) Specify if you require checking of the destination MAC address of the incoming packets on the receiving SFP at the [edit interfaces ge-3/0/0 tdm-options] hierarchy. If you have configured the destination MAC address using the *dmac-address* option, use this option to verify the MAC address on the receiving SFP. If you have enabled MAC address verification and the MAC address does not match, the packet is discarded by the smart SFP.

```
[edit interfaces ge-3/0/0 tdm-options]
user@host # set ces-psn-port-dmac-check-enable
```

8. (Optional) Enable looping back of the input path of TDM traffic on the SFP TDM port. The input path refers to the traffic from the TDM side that is looped back.

```
[edit interfaces ge-3/0/0 tdm-options]
user@host # set tdm-in-loop
```

9. (Optional) Enable looping back of the output path of TDM traffic on the SFP TDM port. The output path refers to the traffic from the Ethernet side that is looped back.

```
[edit interfaces ge-3/0/0 tdm-options]
user@host # set tdm-out-loop
```

Verification

IN THIS SECTION

- [Verifying the STM1 Smart SFP Statistics on the Interface | 139](#)
- [Verifying the STM1 Smart SFP Defects on the Interface | 141](#)

To verify that the STM1 Smart SFP is configured on the router, perform the following tasks:

Verifying the STM1 Smart SFP Statistics on the Interface

Purpose

To verify that the STM1 Smart SFP is configured on the router and to view the STM1 Smart SFP statistics.

Action

To view the STM1 Smart SFP statistics on the Interface, use the `show interfaces ge-3/0/0 smart-sfp-statistics` command.

```
user@host > show interfaces ge-3/0/0 smart-sfp-statistics
Physical interface: ge-3/0/0, Enabled, Physical link is Up
  Interface index: 281, SNMP ifIndex: 742
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
  None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags    : Present Running
```

```

Interface flags: SNMP-Traps Internal: 0x4000
Link flags      : None
Smart Transceiver Type: STM1
Smart SFP: Configurable SFP
Smart SFP Ethernet port[P1] Statistics:          Counters
  Rx frames                  1187126
  Rx bytes                  149855236
  Rx errored fcs frames     0
  Rx unicast frames         1187124
  Rx multicast frames       2
  Rx broadcast frames       0
  Rx fragments               0
  Rx undersize frames       0
  Rx oversize frames        0
  Rx invalid vlan mismatch frames 0
  Tx frames                  1392780998
  Tx bytes                  1796396824
  Tx unicast frames         1377974
  Tx multicast frames        1391403024
  Tx broadcast frames       0
Smart SFP AV IWF Encap/Decap Statistics:          Counters
  Rx Packets                0
  Tx Packets                0
  Malformed Packets          0
  Reordered Packets          0
  Misordered Dropped Packets 0
  Missing Packets            0
  PlayedOut Packets          0
  JitterBuffer Overrun       0
  JitterBuffer Underrun      0
Smart SFP STM1 port[P0] statistics:          Counters
  BiPolarVariations/Excessive zero errors      0
  Tx B3 Errors                  0
  Code Violation path errors      0
Logical interface ge-3/1/0.0 (Index 350) (SNMP ifIndex 615)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Input packets : 1384454023
  Output packets: 0
  Protocol multiservice, MTU: Unlimited

```

Meaning

The STM1 Smart SFP is configured on the router and you can view the STM1 Smart SFP statistics.

Verifying the STM1 Smart SFP Defects on the Interface

Purpose

To verify that the STM1 Smart SFP is configured on the MX480 router and to view the STM1 Smart SFP defects.

Action

To view the STM1 Smart SFP defects on the Interface, use the `show interfaces ge-3/0/0 smart-sfp-defects` command.

```
user@host > show interfaces ge-3/0/0 smart-sfp-defects
Physical interface: ge-3/0/0, Enabled, Physical link is Up
  Interface index: 192, SNMP ifIndex: 546, Generation: 269
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None,
  Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Enabled,
  Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags : None
  Smart Transceiver Type: STM1
  Smart SFP: Configurable Type
  Smart SFP Configurations
    Loopback: None
    DMAC check Enabled: enabled
    Encapsulation Circuit Id: 1000
    Decapsulation Circuit Id: 2000
    VLAN1: 100
    Mode: MEF8
    Dest MAC: 10:0e:7e:37:cd:29
  Smart SFP Defects:
    TDM defects : Local Packet Loss
    System defects : None
```

```

Logical interface ge-3/0/0.0 (Index 369) (SNMP ifIndex 799) (Generation 190)
  Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
  Input packets : 2400
  Output packets: 0
  Protocol multiservice, MTU: Unlimited

```

Meaning

The STM1 Smart SFP is configured on the router and you can view the STM1 Smart SFP defects.

Layer 2 Overhead Attribute in Interface Statistics

SUMMARY

Learn about the Layer 2 overhead attributes, the guidelines for configuring the calculation of layer 2 overhead, and view the layer 2 overhead bytes calculated for ingress and egress traffic on Ethernet Interfaces.

IN THIS SECTION

- [Accounting of the Layer 2 Overhead Attribute in Interface Statistics | 142](#)
- [Configure Layer 2 Overhead Accounting in Interface Statistics | 145](#)
- [Verify the Accounting of Layer 2 Overhead in Interface Statistics | 147](#)

By default, the physical interface and logical interface statistics do not account for Layer 2 overhead in input or output statistics. You can now configure the logical interface statistics to calculate and display all the Layer 2 header details for ingress and egress interfaces.

Accounting of the Layer 2 Overhead Attribute in Interface Statistics

IN THIS SECTION

- [Guidelines to Configure the Computation of Layer 2 Overhead in Interface Statistics | 144](#)

On a router, you can configure the *logical interface* statistics to include the Layer 2 overhead size (header and trailer bytes) for both ingress and egress interfaces. Both the transit and total statistical information are computed and displayed for each logical interface. This functionality is supported on 1-Gigabit, 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces on Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs) on routers.

You can also configure the capability to compute the Layer 2 overhead bytes in interface statistics on Type-3, Type-4 and Type-5 Flexible Port Concentrators (FPCs) on routers. To enable the Layer 2 overhead bytes to be counted in the interface statistics at the PIC level, you must use the `account-layer2-overhead` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level.

If you configure this capability, all the Layer 2 header details (Layer 2 header and cyclic redundancy check [CRC]) based on the Layer 2 encapsulation configured for an interface are calculated and displayed in the logical interface statistics for ingress and egress interfaces in the output of the `show interfaces interface-name` commands. For logical interfaces, the **Input bytes** and **Output bytes** fields under the **Traffic statistics** section in the output of the `show interfaces interface-name <detail | extensive>` command include the Layer 2 overhead of the packets. For logical interfaces, the **Input rate** and **Output rate** fields under the **Traffic statistics** section in the output of the `show interfaces interface-name <media | statistics>` command include the Layer 2 overhead of the packets. For logical interfaces, the values for the newly added **Egress account overhead** and **Ingress account overhead** fields display the Layer 2 overhead size for transmitted and received packets respectively.

The input and output octets at the logical interface configured on the PIC includes all the Layer 2 headers. All the logical interfaces on the PIC, including the ae and the non-ae interfaces, are processed for Layer 2 overhead accounting for the arriving and exiting packets. This method of operation impacts the transit statistics that are primarily used for subscriber accounting and billing purposes in customer networks.

[Table 5 on page 143](#) lists the adjustment bytes that are counted based on the encapsulation on the logical interface over the Ethernet interface, when you enable accounting of Layer 2 overhead in interface statistics at the PIC level. The values for the adjustment bytes that are listed for all types of encapsulation are the same for DPCs and MPCs, with the only exception being for the VLAN CCC adjustment value. On DPCs, the VLAN CCC adjustment value is -4 bytes and on MPCs, the VLAN CCC adjustment value is +4 bytes.

Table 5: Adjustment Bytes for Logical Interfaces over Ethernet Interfaces

Encapsulation Type on Logical Interfaces	Number of Adjustment Bytes	Description
Ethernet DIXv2 (IP datagrams over Ethernet)	18	Untagged (includes CRC)

Table 5: Adjustment Bytes for Logical Interfaces over Ethernet Interfaces (Continued)

Encapsulation Type on Logical Interfaces	Number of Adjustment Bytes	Description
Ethernet DIXv2 (IP datagrams over Ethernet)	22	Single-tagged (includes CRC)
Ethernet DIXv2 (IP datagrams over Ethernet)	26	Double-tagged (includes CRC)
VLAN Bridge	4	CRC
VLAN CCC	4	CRC
VLAN TCC	18	Untagged (includes CRC)
VLAN TCC	22	Single-tagged (includes CRC)
VLAN TCC	26	Double-tagged (includes CRC)
VLAN VPLS	4	CRC

Guidelines to Configure the Computation of Layer 2 Overhead in Interface Statistics

Keep the following points in mind when you configure the computation of Layer 2 overhead in interface statistics:

- When you configure a native VLAN ID on a logical interface, the Layer 2 header adjustment for input statistics is different for tagged and untagged packets. For such interfaces, if you configure the setting to account for Layer 2 overhead, incorrect statistics might be displayed.
- An untagged packet is considered as a tagged packet and an additional 4 bytes are appended to the counter values displayed in the output of the `show interface` command.
- The computed statistics might not be completely accurate in scenarios where the packets are dropped after they have been included in the interface statistics, but before the packets reach the destination.

- Label-switched interface (LSI) statistics on the ingress direction of interfaces do not include the Layer 2 overhead bytes because this functionality of accounting Layer 2 overhead is not supported for such LSI interfaces.
- Layer 2 overhead accounting is not supported for inline service (si) interfaces.
- The total statistics of interfaces do not indicate the complete Layer 2 adjusted statistics. This behavior occurs because the total statistics count is the sum of transit and local statistics. Only the transit statistics are adjusted for Layer 2 and the local statistics are not adjusted for Layer 2.
- Statistics on ae interfaces are calculated in the same manner as non-ae interfaces.
- Adjustment bytes are applicable only for transit statistics that are displayed for logical interfaces.
- For physical interfaces, the adjustment bytes for transit traffic and the non-adjusted bytes for local or protocol-specific traffic are combined and displayed in the output of the `show interfaces` command. (Segregation is not possible.)
- Layer 2 overhead accounting can be enabled at both PIC level and logical interface level.
- When the `account-layer2-overhead` statement is configured, the Layer 2 overhead size in both input and output statistics is accounted for in Dense Port Concentrator (DPCs) and Modular Port Concentrator (MPCs).
- This `account-layer2-overhead` configuration now supports Layer 2 accounting for the Ethernet bridge encapsulation.
- The Layer 2 overhead bytes in interface statistics are saved across a unified ISSU or a *graceful Routing Engine switchover* (GRES) operation.

Configure Layer 2 Overhead Accounting in Interface Statistics

IN THIS SECTION

- [Enable the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level | 146](#)

This topic contains sections that describe the configuration of Layer 2 overhead accounting for interface statistics at the PIC level and logical interface level.

Layer 2 overhead accounting can be enabled at both PIC level and logical interface level through configuration. By default, the physical interface and logical interface statistics do not account for Layer 2 overhead size (header and trailer) in both input and output statistics.

When the `account-layer2-overhead` statement is configured, the Layer 2 overhead size in both input and output statistics is accounted for in the Dense Port Concentrator (DPCs) and the Modular Port Concentrator (MPCs). This `account-layer2-overhead` configuration now supports Layer 2 accounting for the Ethernet bridge encapsulation.

Enable the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level

You can configure the `account-layer2-overhead` statement at the `edit chassis fpc slot-number pic pic-number` hierarchy level to enable accounting of Layer 2 overhead bytes in the ingress and egress interface statistics at the PIC level.



CAUTION: If you modify the setting for accounting of Layer 2 overhead bytes at the PIC level, the PIC is rebooted, causing all of the physical and logical interfaces to be deleted and readded on the PIC. Due to this behavior, we recommend that you exercise caution while using this feature.

The computation method of Layer 2 overhead on different interface types is as follows:

- For Ethernet interfaces, all the Layer 2 headers are counted.
- For non-Ethernet interfaces, the Frame Relay, PPP, or Cisco HDLC headers are counted, while the bit or byte stuffing headers are excluded.

To enable accounting of Layer 2 overhead at the PIC level for ingress and egress traffic on interfaces:

1. Access a DPC or an MPC-occupied slot and the PIC where the interface is to be enabled.

```
[edit chassis]
user@host# edit fpc slot-number pic number
```

2. Specify the Layer 2 overhead value in bytes that is the octet adjustment per packet added to the total octet count for ingress and egress traffic on all the interfaces in the PIC.

```
[edit chassis fpc slot-number pic number]
user@host# set account-layer2-overhead
```

Verify the Accounting of Layer 2 Overhead in Interface Statistics

IN THIS SECTION

- [Purpose | 147](#)
- [Action | 147](#)

Purpose

Display information about the Layer 2 overhead bytes that are counted in interface statistics for egress and ingress traffic on Ethernet interfaces.

Action

- To display information about the Layer 2 overhead bytes that are counted in interface statistics:

For physical and logical interfaces, the values displayed for the Input rate and Output rate fields under the Traffic statistics section include the Layer 2 overhead of the packets.

```
user@host> show interfaces ge-5/2/0 statistics detail
```

```
Physical interface: ge-5/2/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 519, Generation: 149
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, 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: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:61:d9:74, Hardware address: 00:1d:b5:61:d9:74
  Last flapped    : 2009-11-11 11:24:00 PST (09:23:08 ago)
  Statistics last cleared: 2009-11-11 17:50:58 PST (02:56:10 ago)
  Traffic statistics:
```

```

Input bytes : 271524 0 bps
Output bytes : 37769598 352 bps
Input packets: 3664 0 pps
Output packets: 885790 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 16681118
Input packets: 0
Output packets: 362633
Multicast statistics:
IPV4 multicast statistics:
Input bytes : 112048 0 bps
Output bytes : 20779920 0 bps
Input packets: 1801 0 pps
Output packets: 519498 0 pps
IPV6 multicast statistics:
Input bytes : 156500 0 bps
Output bytes : 16681118 0 bps
Input packets: 1818 0 pps
Output packets: 362633
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0,
L2 channel errors: 0,
L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0,
Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
0 best-effort 882558 882558 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 3232 3232 0
Active alarms : None
Active defects : None

Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
Input bytes : 271524

```

```

Output bytes : 37769598
Input packets: 3664
Output packets: 885790
IPv6 transit statistics:
Input bytes : 0
Output bytes : 16681118
Input packets: 0
Output packets: 362633
Local statistics:
Input bytes : 271524
Output bytes : 308560
Input packets: 3664
Output packets: 3659
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 37461038 0 bps
Input packets: 0 0 pps
Output packets: 882131 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 16681118
Input packets: 0
Output packets: 362633
Multicast statistics:
IPV4 multicast statistics:
Input bytes : 112048 0 bps
Output bytes : 20779920 0 bps
Input packets: 1801 0 pps
Output packets: 519498 0 pps
IPV6 multicast statistics:
Input bytes : 156500 0 bps
Output bytes : 16681118 0 bps
Input packets: 1818 0 pps
Output packets: 362633 0 pps
Protocol inet, MTU: 1500, Generation: 151, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 40.40.40.0/30, Local: 40.40.40.2, Broadcast: 40.40.40.3, Generation: 167
Protocol inet6, MTU: 1500, Generation: 152, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: ::40.40.40.0/126, Local: ::40.40.40.2
Generation: 169
Addresses, Flags: Is-Preferred
Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:d974

```

```
Protocol multiservice, MTU: Unlimited, Generation: 171
Generation: 153, Route table: 0
Policer: Input: __default_arp_policer__
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
13.3	Starting with Junos OS Release 13.3, account-layer2-overhead is not supported on MX Series routers with MPC3E (on both PIC and logical interface levels).
13.2	Starting with Junos OS Release 13.2, configuring the logical interface statistics to include Layer 2 is supported on 10-Gigabit Ethernet interfaces on MX Series routers with MPC4E.

RELATED DOCUMENTATION

[Ethernet Interfaces Overview | 2](#)

[Configure Ethernet Interfaces | 4](#)

[100GbE, 40GbE, and 10GbE | 351](#)

Gigabit Ethernet Policers

SUMMARY

Policers enable you to perform simple traffic policing on Gigabit Ethernet Interfaces without configuring a firewall filter. You can use this topic to configure an input priority map, an output priority map, and then apply the policy. Use this topic for information about how to configure a two-color policer and tri-color policer.

IN THIS SECTION

- [Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 151](#)
- [Configure Gigabit Ethernet Policers | 153](#)
- [Configure Gigabit Ethernet Two-Color and Tricolor Policers | 159](#)

Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs

For Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs, you can configure granular per-VLAN class-of-service (CoS) capabilities and extensive instrumentation and diagnostics on a per-VLAN and per-MAC address basis.

VLAN rewrite, tagging, and deleting enables you to use VLAN address space to support more customers and services.

VPLS allows you to provide a point-to-multipoint LAN between a set of sites in a VPN. Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) are combined with VPLS to deliver metro Ethernet service.

For Gigabit Ethernet IQ2 and IQ2-E and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, you can apply Layer 2 policing to logical interfaces in the egress or ingress direction. Layer 2 policers are configured at the [edit firewall] hierarchy level. You can also control the rate of traffic sent or received on an interface by configuring a policer overhead at the [edit chassis fpc *slot-number* pic *slot-number*] hierarchy level.

[Table 6 on page 151](#) lists the capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Table 6: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
Layer 2		
802.3ad link aggregation	Yes	Yes
Maximum VLANs per port	384	1023
Maximum transmission unit (MTU) size	9192	9192
MAC learning	Yes	Yes
MAC accounting	Yes	Yes
MAC filtering	Yes	Yes

Table 6: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs (Continued)

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
Destinations per port	960	960
Sources per port	64	64
Hierarchical MAC policers	Yes, premium and aggregate	No, aggregate only
Multiple TPID support and IP service for nonstandard TPIDs	Yes	Yes
Multiple Ethernet encapsulations	Yes	Yes
Dual VLAN tags	Yes	No
VLAN rewrite	Yes	No

Layer 2 VPNs

VLAN CCC	Yes	Yes
Port-based CCC	Yes	Yes
Extended VLAN CCC Virtual Metropolitan Area Network (VMAN) Tag Protocol	Yes	Yes

CoS

PIC-based egress queues	Yes	Yes
Queued VLANs	Yes	No
VPLS	Yes	Yes

For more information about configuring VPLS, see the [Junos OS VPNs Library for Routing Devices](#).

You can also configure CoS on logical IQ interfaces. For more information, see the [Junos OS Class of Service User Guide for Routing Devices](#).

SEE ALSO

[MAC Address Accounting for Ethernet Interfaces](#) | 9

Configure Gigabit Ethernet Policers

IN THIS SECTION

- [Overview](#) | 153
- [Configure a Policer](#) | 154
- [Specify an Input Priority Map](#) | 155
- [Specify an Output Priority Map](#) | 155
- [Apply a Policer](#) | 156
- [Example: Configure Gigabit Ethernet Policers](#) | 157

Overview

On Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs, you can define rate limits for premium and aggregate traffic received on the interface. These policers allow you to perform simple traffic policing without configuring a firewall filter. First you configure the Ethernet policer profile, next you classify ingress and egress traffic, then you can apply the policer to a logical interface.

For Gigabit Ethernet PICs with SFPs, the policer rates you configure can be different than the rates on the Packet Forward Engine. The difference results from Layer 2 overhead. The PIC accounts for this difference.

On routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.

- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

Configure a Policer

To configure an Ethernet policer profile, include the `ethernet-policer-profile` statement at the `[edit interfaces interface-name gigether-options ethernet-switch-profile]` hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
ethernet-policer-profile {
    policer cos-policer-name {
        aggregate {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
        premium {
            bandwidth-limit bps;
            burst-size-limit bytes;
        }
    }
}
```

In the Ethernet policer profile, the aggregate-priority policer is mandatory; the premium-priority policer is optional.

For aggregate and premium policers, you specify the bandwidth limit in bits per second. You can specify the value 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). There is no absolute minimum value for bandwidth limit, but any value below 61,040 bps will result in an effective rate of 30,520 bps. The maximum bandwidth limit is 4.29 Gbps.

The maximum burst size controls the amount of traffic bursting allowed. To determine the burst-size limit, you can multiply the bandwidth of the interface on which you are applying the filter by the amount of time you allow a burst of traffic at that bandwidth to occur:

$$\text{burst size} = \text{bandwidth} \times \text{allowable time for burst traffic}$$

If you do not know the interface bandwidth, you can multiply the maximum MTU of the traffic on the interface by 10 to obtain a value. For example, the burst size for an MTU of 4700 would be 47,000 bytes. The burst size should be at least 10 interface MTUs. The maximum value for the burst-size limit is 100 MB.

Specify an Input Priority Map

An input priority map identifies ingress traffic with specified IEEE 802.1p priority values, and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an input priority map by including the `ieee802.1p premium [values]` statement at the `[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]` hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]
ieee802.1p premium [ values ];
```

The priority values can be from 0 through 7. The remaining traffic is classified as nonpremium (or aggregate).

Specify an Output Priority Map

An output priority map identifies egress traffic with specified queue classification and packet loss priority (PLP), and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an output priority map by including the `classifier` statement at the `[edit interfaces interface-name gigether-options ethernet-policer-profile output-priority-map]` hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile output-priority-map]
classifier {
    premium {
        forwarding-class class-name {
            loss-priority (high | low);
        }
    }
}
```

You can define a forwarding class, or you can use a predefined forwarding class. [Table 7 on page 155](#) shows the predefined forwarding classes and their associated queue assignments.

Table 7: Default Forwarding Classes

Forwarding Class Name	Queue
best-effort	Queue 0

Table 7: Default Forwarding Classes (Continued)

Forwarding Class Name	Queue
expedited-forwarding	Queue 1
assured-forwarding	Queue 2
network-control	Queue 3

Apply a Policer

On all Ethernet interfaces, Gigabit Ethernet IQ, IQ2, and IQ2-E PICs, and Gigabit Ethernet PICs with SFPs, you can apply input and output policers that define rate limits for premium and aggregate traffic received on the logical interface. Aggregate policers are supported on Gigabit Ethernet PICs with SFPs.

These policers allow you to perform simple traffic policing without configuring a firewall filter.

To apply policers to specific source MAC addresses, include the `accept-source-mac` statement:

```
accept-source-mac {
    mac-address mac-address {
        policer {
            input cos-policer-name;
            output cos-policer-name;
        }
    }
}
```

You can include these statements at the following hierarchy levels:

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

You can specify the MAC address as *nn.nn.nn.nn.nn.nn* or *nnnn.nnnn.nnnn*, where *n* is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include multiple `mac-address` statements in the logical interface configuration.

On untagged Gigabit Ethernet interfaces you should not configure the `source-address-filter` statement at the [edit interfaces *ge-fpc/pic/port* *giether-options*] hierarchy level and the `accept-source-mac` statement at the [edit interfaces *ge-fpc/pic/port* *giether-options* unit *logical-unit-number*] hierarchy level simultaneously.

If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces you should not configure the source-address-filter statement at the [edit interfaces ge-fpc/pic/port gigether-options] hierarchy level and the accept-source-mac statement at the [edit interfaces ge-fpc/pic/port gigether-options unit *logical-unit-number*] hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

If the remote Ethernet card is changed, the interface does not accept traffic from the new card because the new card has a different MAC address.

The MAC addresses you include in the configuration are entered into the router's MAC database. To view the router's MAC database, enter the `show interfaces mac-database interface-name` command:

```
user@host> show interfaces mac-database interface-name
```

In the input statement, list the name of one policer template to be evaluated when packets are received on the interface. In the output statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.

On IQ2 and IQ2-E PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

You can use the same policer one or more times. If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

Example: Configure Gigabit Ethernet Policers

IN THIS SECTION

- [Example | 157](#)
- [Example Configuration | 158](#)

Example

This example illustrates the following:

- Configure interface ge-6/0/0 to treat priority values 2 and 3 as premium. On ingress, this means that IEEE 802.1p priority values 2 and 3 are treated as premium. On egress, it means traffic that is classified into queue 0 or 1 with PLP of low and queue 2 or 3 with PLP of high, is treated as premium.
- Define a policer that limits the premium bandwidth to 100 Mbps and burst size to 3 k, and the aggregate bandwidth to 200 Mbps and burst size to 3 k.
- Specify that frames received from the MAC address 00:01:02:03:04:05 and the VLAN ID 600 are subject to the policer on input and output. On input, this means frames received with the source MAC address 00:01:02:03:04:05 and the VLAN ID 600 are subject to the policer. On output, this means frames transmitted from the router with the destination MAC address 00:01:02:03:04:05 and the VLAN ID 600 are subject to the policer.

Example Configuration

```
[edit interfaces]
ge-6/0/0 {
    gigether-options {
        ether-switch-profile {
            ether-policer-profile {
                input-priority-map {
                    ieee-802.1p {
                        premium [ 2 3 ];
                    }
                }
                output-priority-map {
                    classifier {
                        premium {
                            forwarding-class best-effort {
                                loss-priority low;
                            }
                            forwarding-class expedited-forwarding {
                                loss-priority low;
                            }
                            forwarding-class assured-forwarding {
                                loss-priority high;
                            }
                            forwarding-class network-control {
                                loss-priority high;
                            }
                        }
                    }
                }
            }
        }
    }
}
```

```
        }
        policer policer-1 {
            premium {
                bandwidth-limit 100m;
                burst-size-limit 3k;
            }
            aggregate {
                bandwidth-limit 200m;
                burst-size-limit 3k;
            }
        }
    }
    unit 0 {
        accept-source-mac {
            mac-address 00:01:02:03:04:05 {
                policer {
                    input policer-1;
                    output policer-1;
                }
            }
        }
    }
}
```

Configure Gigabit Ethernet Two-Color and Tricolor Policers

IN THIS SECTION

- [Overview | 160](#)
- [Configure a Policer | 161](#)
- [Apply a Policer | 162](#)
- [Configure and Apply a Policer | 162](#)

Overview

For Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers, you can configure two-color and tricolor marking policers and apply them to logical interfaces to prevent traffic on the interface from consuming bandwidth inappropriately.

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

Juniper Networks router architectures support three types of policer:

- Two-color policer—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit in some way, or simply discard them. A policer is most useful for metering traffic at the port (physical interface) level.
- Single-rate tricolor marking (single-rate TCM)—A single-rate tricolor marking policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).

Starting in Junos OS Release 13.1, traffic is classified into three categories: Green, Red, and Yellow. Following list describes the categories:

- Green—Burst size of the packets that arrive is less than the sum of the configured CIR and CBS.
- Red—Burst size of the packets that arrive is greater than the sum of the configured CIR and EBS.
- Yellow—Burst size of the packets that arrive is greater than the CBS but less than the EBS.

Single-rate TCM is most useful when a service is structured according to packet length and not peak arrival rate.

- Two-rate Tricolor Marking (two-rate TCM)—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and peak information rate (PIR), along with their associated burst sizes, the CBS and EBS.

Traffic is classified into the following three categories:

- Green—Burst size of the packets that arrive is less than the sum of the configured CIR and CBS.
- Red—Burst size of the packets that arrive is greater than the sum of the configured PIR and EBS.
- Yellow—Traffic does not belong to either the green or the red category.

Two-rate TCM is most useful when a service is structured according to arrival rates and not necessarily packet length.



NOTE: Unlike policing (described in "Configuring Gigabit Ethernet Policers" on page 153), configuring two-color policers and tricolor marking policers requires that you configure a firewall filter.

Configure a Policer

Two-color and tricolor marking policers are configured at the [edit firewall] hierarchy level.

A tricolor marking policer polices traffic on the basis of metering rates, including the CIR, the PIR, their associated burst sizes, and any policing actions configured for the traffic.

To configure tricolor policer marking, include the three-color-policer statement with options at the [edit firewall] hierarchy level:

```
[edit firewall]
three-color-policer name {
    action {
        loss-priority high {
            then discard;
        }
    }
    single-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-information-rate bps;
        committed-burst-size bytes;
        peak-information-rate bps;
        peak-burst-size bytes;
    }
}
```

```

    }
}
```

For more information about configuring tricolor policer markings, see the [Routing Policies, Firewall Filters, and Traffic Policers User Guide](#) and the [Junos OS Class of Service User Guide for Routing Devices](#).

Apply a Policer

Apply a two-color policer or tricolor policer to a logical interface to prevent traffic on the interface from consuming bandwidth inappropriately. To apply two-color or tricolor policers, include the `layer2-policer` statement:

```

layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    policer-name;
}
```

You can include these statements at the following hierarchy levels:

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

Use the `input-policer` statement to apply a two-color policer to received packets on a logical interface and the `input-three-color` statement to apply a tricolor policer. Use the `output-policer` statement to apply a two-color policer to transmitted packets on a logical interface and the `output-three-color` statement to apply a tricolor policer. The specified policers must be configured at the [edit firewall] hierarchy level. For each interface, you can configure a three-color policer or two-color input policer or output policers—you cannot configure both a three-color policer and a two-color policer.

Configure and Apply a Policer

Configure tricolor policers and apply them to an interface:

```

[edit firewall]
three-color-policer three-color-policer-color-blind {
    logical-interface-policer;
    two-rate {
        color-blind;
        committed-information-rate 1500000;
```

```

        committed-burst-size 150;
        peak-information-rate 3;
        peak-burst-size 300;
    }
}

three-color-policer three-color-policer-color-aware {
    logical-interface-policer;
    two-rate {
        color-aware;
        committed-information-rate 1500000;
        committed-burst-size 150;
        peak-information-rate 3;
        peak-burst-size 300;
    }
}
[edit interfaces ge-1/1/0]
unit 1 {
    layer2-policer {
        input-three-color three-color-policer-color-blind;
        output-three-color three-color-policer-color-aware;
    }
}

```

Configure a two-color policer and apply it to an interface:

```

[edit firewall]
policer two-color-policer {
    logical-interface-policer;
    if-exceeding {
        bandwidth-percent 90;
        burst-size-limit 300;
    }
    then loss-priority-high;
}
[edit interfaces ge-1/1/0]
unit 2 {
    layer2-policer {
        input-policer two-color-policer;
        output-policer two-color-policer;
    }
}

```

RELATED DOCUMENTATION

MAC Address Accounting for Ethernet Interfaces | [9](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
13.1	Starting in Junos OS Release 13.1, traffic is classified into three categories: Green, Red, and Yellow.

Gigabit Ethernet Autonegotiation

SUMMARY

Learn about Autonegotiation and how to configure Autonegotiation on Gigabit Ethernet Interfaces.

IN THIS SECTION

- [Gigabit Ethernet Autonegotiation Overview | 164](#)
- [Configure Gigabit Ethernet Autonegotiation | 165](#)

Gigabit Ethernet Autonegotiation Overview

Autonegotiation is enabled by default on all Gigabit Ethernet and Tri-Rate Ethernet copper interfaces. However, you can explicitly enable autonegotiation to configure remote fault options manually. When the autonegotiation is disabled, the speed has to be explicitly configured to 10–100 Mbps.

When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.

Configure Gigabit Ethernet Autonegotiation

IN THIS SECTION

- [Configure Gigabit Ethernet Autonegotiation with Remote Fault | 165](#)
- [Configure Autonegotiation Speed | 165](#)
- [Autonegotiation Status Display | 166](#)

Configure Gigabit Ethernet Autonegotiation with Remote Fault

To configure explicit autonegotiation and remote fault, include the `auto-negotiation` statement and the `remote-fault` option at the `[edit interfaces ge-fpc/pic/port gigether-options]` hierarchy level.

```
[edit interfaces ge-fpc/pic/port gigether-options]
(auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online | local-interface-
offline>
```

Configure Autonegotiation Speed

Routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs support autonegotiation of speed. The autonegotiation specified interface speed is propagated to CoS, routing protocols, and other system components. Half-duplex mode is not supported.

Routers with IQ2 PICs connected to other devices require matching auto-negotiation configurations for both the PIC and for the device in order to achieve link up.

To specify the autonegotiation speed, use the `speed` (`auto` | `1Gbps` | `100Mbps` | `10Mbps` | `auto-10m-100m`) statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level.

To set port speed negotiation to a specific rate, set the port speed to `1Gbps`, `100Mbps`, or `10Mbps`. If the negotiated speed and the interface speed do not match, the link will not be brought up.

If you set the autonegotiation speed `auto` option, then the port speed is negotiated.

Starting from Junos OS Release 14.2, the `auto-10m-100m` option allows the fixed tri-speed port to auto negotiate with ports limited by `100m` or `10m` maximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC.

You can disable auto MDI/MDIX using the `no-auto-mdix` statement at the `[edit interfaces ge-fpc/pic/port gigether-options]` hierarchy level.

Use the `show interfaces ge- fpc/pic/port brief` command to display the auto negotiation of speed and auto MDI/MDIX states.

Starting in Junos OS Release 14.2, on routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the `auto-10m-100m` option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link. However, when you view the interface speed of the DPC, using the `show interfaces` command, the value of the speed is not accurately displayed. For instance, if you configure the speed of the Tri-rate enhanced DPC, as 100Mbps on both sides of the link, the interface speed of the DPC is negotiated to 100 Mbps. However, the interface speed of the DPC displays 1 bps. This is an issue with the `show interfaces` command only. The actual interface speed is 100 Mbps.

Autonegotiation Status Display

To display Gigabit Ethernet interface details, including the autonegotiation status, use the operational mode command `show interfaces ge- fpc/pic/port extensive`.

[Table 8 on page 166](#) and [Table 9 on page 170](#) provide information about the autonegotiation status on local and remote routers with fiber interfaces. The status of the link and LED can vary depending on the level of autonegotiation set and the transmit and receive fiber status.

Table 8: Mode and Autonegotiation Status (Local)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Red	DOWN	
ON	ON	Default	Green	UP	No-autonegotiation
ON	OFF	Default	Red	DOWN	

Table 8: Mode and Autonegotiation Status (Local) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Green	UP	
ON	ON	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	OFF	No-autonegotiation	Red	DOWN	
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit	Green	UP	No-autonegotiation
ON	OFF	Explicit	Red	DOWN	

Table 8: Mode and Autonegotiation Status (Local) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
OFF	ON	Explicit	Green	UP	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Green	UP	Complete
OFF	ON	Explicit+RFI-Offline	Red	DOWN	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Green	UP	No-autonegotiation
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Offline	Green	UP	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	

Table 8: Mode and Autonegotiation Status (Local) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	No-autonegotiation*
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Green	UP	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	Complete

Table 9: Mode and Autonegotiation Status (Remote)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete
ON	ON	Default	Red	DOWN	
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	ON	No-autonegotiation	Red	DOWN	
ON	OFF	No-autonegotiation	Red	DOWN	
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	ON	Explicit	Red	DOWN	
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	

Table 9: Mode and Autonegotiation Status (Remote) *(Continued)*

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Explicit	Red	DOWN	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit+RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit+RFI-Offline	Red	DOWN	
OFF	ON	Explicit+RFI-Offline	Red	DOWN	
OFF	OFF	Explicit+RFI-Offline	Red	DOWN	
ON	ON	Explicit+RFI-Online	Green	UP	Complete
ON	ON	Explicit+RFI-Online	Red	DOWN	
ON	OFF	Explicit+RFI-Online	Red	DOWN	
OFF	ON	Explicit+RFI-Online	Red	DOWN	
OFF	OFF	Explicit+RFI-Online	Red	DOWN	

4

PART

Port Speed

- Port Speed Overview | **173**
- Configure Port Speed | **176**
- Port Speed on ACX Routers | **182**
- Port Speed on MX Routers | **210**
- Port Speed on PTX Routers | **274**

Port Speed Overview

IN THIS SECTION

- [Introduction to Port Speed | 173](#)
- [Platform-Specific Port Profile Behavior | 175](#)

Learn about port speed on a device or line card, support for multiple port speed details, guidelines and how to configure the port speed.

Introduction to Port Speed

IN THIS SECTION

- [Interface Naming Conventions | 174](#)
- [What is Oversubscription? | 175](#)

Port speed refers to the maximum amount of data that the line card transmits through a port at any given second. Port speed is measured as follows:

- Kilobits per second (Kbps)
- Gigabits per second (Gbps)
- Terabits per second (Tbps)

[Table 10 on page 174](#) describes the different types of port speed configuration.

[Table 11 on page 175](#) describes the naming formats for the channelized and non-channelized interfaces.

["Configure Port Speed - PIC Level" on page 177](#) describes the port speed configuration at PIC Level - Chassis Hierarchy.

"Configure Port Speed - Port Level" on page 178 describes the port speed configuration at Port Level - Chassis Hierarchy.

"Configure Port Speed for Non-Channelized Interfaces" on page 179 describes the steps to configure the port speed for non-channelized interfaces from the [edit interfaces] hierarchy.

"Configure Port Speed for Channelized Interfaces" on page 180 describes the steps to configure the port speed for channelized interfaces from the [edit interfaces] hierarchy.

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Review the "Platform-Specific Port Profile Behavior" on page 175 section for notes related to your platform.

Table 10: Types of Port Speed Configuration

Port Speed Configuration Type	Description
PIC or MIC Level	You can configure all the ports in one PIC or MIC to operate at the same speed. For example, you can configure all the ports of a PIC that support port speed of 100 Gbps to operate at 100 Gbps speed. If you do not specify the speed by using the <code>pic-mode</code> statement, then the port operates in the default speed.
Port Level	You can configure each port to operate at a different speed and thus enable each port. When you configure the port speed at the port level, you have the flexibility of operating the ports of the line card at different supported speeds. When you change the speed of a specific port in a given PIC using the <code>speed</code> statement, then only the speed of that port is modified. All other ports in the PIC remain unaffected. For example, you can configure four 10-Gigabit Ethernet interfaces on port 0, one 40-Gigabit Ethernet interface on port 1, and one 100-Gigabit Ethernet interface on port 2.

Interface Naming Conventions

Each interface name includes a unique identifier and follows a naming convention. When you configure the interface, use the interface name. You can either configure a port as a single interface (non-channelized interface) or partition the port into smaller data channels or multiple interfaces (channelized interfaces).

When multiple interfaces are supported on a physical port, you use the colon (:) notation in the interface naming conventions as a delimiter to differentiate the multiple interfaces on a physical port. In the interface naming convention, `xe-x/y/z:channel`:

- x refers to the FPC slot number.
- y refers to the PIC slot number.
- z refers to the physical port number.
- channel refers to the number of channelized interfaces.

When the 40-Gigabit Ethernet interfaces (et-fpc/pic/port) are channelized as 10-Gigabit Ethernet interfaces, the interface appears in the xe-fpc/pic/port:channel format, and channel is a value of 0 through 3.

Table 11: Channelized and Non-Channelized Interface Naming Formats

Interfaces	Non-channelized Interfaces Naming Formats	Channelized Interfaces Naming Formats
10-Gigabit Ethernet Interfaces	Prefix is xe-. The interface name appears in the xe-fpc/pic/port format.	Prefix is xe-. The interface name appears in the xe-fpc/pic/port:channel format.
25-Gigabit Ethernet Interfaces, 40-Gigabit Ethernet Interfaces, 100-Gigabit Ethernet Interfaces, 200-Gigabit Ethernet Interfaces, and 400-Gigabit Ethernet Interfaces.	Prefix is et-. The interface name appears in the et-fpc/pic/port format.	Prefix is et-. The interface name appears in the et-fpc/pic/port:channel format.

What is Oversubscription?

Oversubscription occurs when you configure the speed of a port at the PIC level, and all ports that support that speed are enabled. To prevent oversubscription, you can configure the number of active ports that operate at the configured speed. Interfaces are created only for active ports. When oversubscription of Packet Forwarding Capacity is not supported, the demand on each Packet Forwarding Engine should be less than or equal to its forwarding capacity.

Platform-Specific Port Profile Behavior

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Use the following table to review platform-specific behavior for your platform:

Platform	Difference
ACX Series	<ul style="list-style-type: none"> ACX710, ACX7100, and ACX7509 routers that support port profile supports channelized and non-channelized configuration at the interface hierarchy. ACX5448 and ACX7024 routers that support port profile supports port level and pic level configuration at the chassis hierarchy.
MX Series	<ul style="list-style-type: none"> MX routers that support port profile supports port level and pic level configuration at the chassis hierarchy.
PTX Series	<ul style="list-style-type: none"> PTX routers that support port profile supports channelized and non-channelized configuration at the interface hierarchy.

Configure Port Speed

IN THIS SECTION

- Configure Port Speed - PIC Level | [177](#)
- Configure Port Speed - Port Level | [178](#)
- Configure Port Speed for Non-Channelized Interfaces | [179](#)
- Configure Port Speed for Channelized Interfaces | [180](#)

This topic provides you with information about how to configure port speed in port level and PIC level at chassis and interface hierarchy.

Configure Port Speed - PIC Level

To configure the operating port speed in the PIC level at chassis hierarchy, follow these steps:

1. In the [edit chassis] *fpc fpc-slot pic pic-number* hierarchy level, configure the port speed using the **set pic-mode *pic-speed*** command.

```
[edit chassis fpc 0 pic 0]
user@host# set pic-mode 100g
```

2. Configure the number of active physical ports that operate at the configured speed.

```
[edit chassis fpc 0 pic 0]
user@host# set number-of-ports 4
```

3. Specify the number of logical interfaces you want to create on a physical port.

```
[edit chassis fpc 0 pic 1]
user@host# set number-of-sub-ports 4
```

4. Specify the port number that you want to power off.

```
[edit chassis fpc 0 pic 1]
user@host# set port 2 unused
```

5. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
pic-mode 100G;
number-of-ports 4;
```

6. Commit the configuration.

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

Configure Port Speed - Port Level

To configure the operating port speed in the port level at chassis hierarchy, follow these steps:

1. In the [edit chassis] fpc *fpc-slot* pic *pic-number* hierarchy level, configure the port speed using the **set port *port-number* speed (10g | 40g | 100g)** command.

```
[edit chassis fpc 0 pic 0]
user@host# set port 0 speed 40g
set port 0 speed 100g
```

2. Specify the number of logical interfaces you want to create on a physical port.

```
[edit chassis fpc 0 pic 1]
user@host# set number-of-sub-ports 4
```

3. Specify the port number that you want to power off.

```
[edit chassis fpc 0 pic 1]
user@host# set port 2 unused
```

4. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show port 0 {
    speed 40g;
}
port 1 {
    speed 100g;
}
[edit chassis fpc 0 pic 1]
user@host# show
port 1 {
    number-of-sub-ports 4;
}
```

5. Commit the configuration.

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

Configure Port Speed for Non-Channelized Interfaces

To configure the operating port speed for non-channelized interface at the interface hierarchy, follow these steps:

1. To specify the port speed, in the [edit interface *interface-name*] hierarchy level, configure the speed value using the set speed (10g | 25g | 40g | 50g | 100g | 400g) command.

```
[edit interfaces et-1/0/3]
user@host# set speed 100g
```

2. To configure the speed for a group of ports, in the [edit wildcard range interface *interface-name*] hierarchy level, configure the speed value using the set speed (10g | 25g | 40g | 50g | 100g | 400g) command.

```
[edit wildcard range set interfaces et-1/0/[0-5]
user@host# set speed 100g
```

3. To control the number of interfaces created on a physical port, use the unused statement.

```
[edit interfaces et-2/0/3]
user@host# set unused
```



NOTE: In this example, no interfaces (channelized or non-channelized) are created on port 3 of the line card installed in the FPC slot 2.

4. Verify the configuration.

```
[edit interfaces]
user@host# et-x/y/z {
    speed 100g;
```

```

unit 0 {
    ...
}
...
unit N {
    ...
}
...
et-x/y/z {
    unused;
}

```

5. Commit the configuration.

```

[edit interfaces et-1/0/3]
user@host# commit

```

Configure Port Speed for Channelized Interfaces

To configure the operating port speed for channelized interface at the interface hierarchy, follow these steps:

1. To specify the port speed, in the [edit interface *interface-name*] hierarchy level, configure the speed value using the **set speed (10g | 25g | 40g | 50g | 100g | 400g)** command.

```

[edit interfaces et-1/0/3]
user@host# set speed 100g

```

2. To configure the speed for a group of ports, in the [edit wildcard range interface *interface-name*] hierarchy level, configure the speed value using the **set speed (10g | 25g | 40g | 50g | 100g | 400g)** command.

```

[edit wildcard range set interfaces et-1/0/[7-12]
user@host# set speed 100g

```

3. To specify the number of interfaces you want to configure per port.

```
[edit interfaces et-2/0/3]
user@host# set number-of-sub-ports 4
```



NOTE: In this example, in Step 1 and Step 2, you configure 4x100GE channelized interfaces.

4. To control the number of interfaces created on a physical port, use the `unused` statement.

```
[edit interfaces et-2/0/4]
user@host# set unused
```

5. Verify the configuration.

```
[edit interfaces]
user@host# et-x/y/z {
    speed 100g;
    number-of-sub-ports 4;
    et-x/y/z:0 {
        unit 0 {
            ...
        }
    }
    et-x/y/z:1 {
        unit 0 {
            ...
        }
    }
    et-x/y/z:2 {
        unit 0 {
            ...
        }
    }
    et-x/y/z:3 {
        unit 0 {
            ...
        }
    }
}
```

```
...
et-x/y/z:6 {
    unused;
}
```

6. Commit the configuration.

```
[edit interfaces et-1/0/3]
user@host# commit
```

Port Speed on ACX Routers

IN THIS SECTION

- [Port Speed on ACX710 Router | 182](#)
- [Port Speed on ACX5448 Router | 184](#)
- [Port Speed on ACX7020 | 187](#)
- [Port Speed on ACX7024 and ACX7024X Router | 188](#)
- [Port Speed on ACX7100 Router | 190](#)
- [Port Speed on ACX7332 and ACX7348 Routers | 197](#)
- [Port Speed on ACX7509 Router | 202](#)

In this topic we have information related to the port speed on a ACX routers and line cards, support for multiple port speed details, guidelines and how to configure the port speed.

Port Speed on ACX710 Router

For information on ACX710 routers see *ACX710 Universal Metro Router Hardware Guide*.

To view the supported transceivers, optical interfaces, and DAC cables on ACX710, see [Hardware Compatibility Tool \(HCT\)](#).

Table 12: Port Speed for ACX710

PIC	Port Number	Port Speed Supported
PIC 0	0 -23	10-Gigabit Ethernet
		1-Gigabit Ethernet
	0 -15	1-Gigabit Ethernet
	16 -23	100-Mbps and 1-Gigabit Ethernet
PIC 1	0 -3	40-Gigabit Ethernet
		4x10-Gigabit Ethernet
		4x25-Gigabit Ethernet
		100-Gigabit Ethernet
		NOTE: By default, all the active ports operate in 100-Gigabit Ethernet mode.

Use the speed command to set the speed on tri-rate copper SFP port.

Table 13: Interface Naming Convention for ACX710

PIC	Interface type	Interfaces
PIC 0	1-Gigabit/10-Gigabit Ethernet interface (24 SFP+ or SFP ports)	xe-0/0/0 - xe-0/0/23
	1-Gigabit Ethernet interface (Tri-rate SFP-T optics)	xe-0/0/0 - xe-0/0/15
	100-Mbps and 1-Gigabit Ethernet interface (Tri-rate SFP-T optics)	xe-0/0/16 - xe-0/0/23

Table 13: Interface Naming Convention for ACX710 (Continued)

PIC	Interface type	Interfaces
PIC 1	100-Gigabit Ethernet interface (4 QSFP28 ports)	et-0/1/0 - et-0/1/3

Port Speed on ACX5448 Router

The ACX5448-D router, supports two ot interfaces (ot-0/2/0 and ot-0/2/1). For each CFP2-DCO optical module installed on ports 38 and 39, one optical transport interface (-ot) is created. You can map two 100-Gigabit Ethernet (et) interfaces can be mapped to each ot interface, depending on the rate (100 Gbps or 200 Gbps) that you configure for the CFP2 ports. As a result, four et interfaces are possible: et-0/2/0, et-0/2/1, et-0/2/2, and et-0/2/3. For information on the interface mapping and modulation format for ACX5448-D routers, see [Interface Mapping and Modulation Format for ACX5448-D Routers](#).

For information on ACX5448 routers see [ACX5448, ACX5448-D, and ACX5448-M Universal Metro Routers Hardware Guide](#).

To view the supported transceivers, optical interfaces, and DAC cables on ACX5448, see [Hardware Compatibility Tool \(HCT\)](#).

Table 14: Port Speed for ACX5448

PIC	Port Number	Port Speed Supported
PIC 0	0 through 47	10-Gigabit Ethernet with SFP+ optics 1-Gigabit Ethernet with SFP optics 100-Megabit Ethernet with SFP optics

Table 14: Port Speed for ACX5448 (Continued)

PIC	Port Number	Port Speed Supported
PIC 1	0 through 3	100-Gigabit Ethernet with QSFP28 optics 4x25-Gigabit Ethernet (with breakout cables and channelization). 40-Gigabit Ethernet with QSFP+ optics 4x10-Gigabit Ethernet (with breakout cables and channelization).

Table 15: Port Speed for ACX5448-D

PIC	Port Number	Port Speed Supported
PIC 0	0 through 35	10-Gigabit Ethernet with SFP+ optics 1-Gigabit Ethernet with SFP optics 100-Megabit Ethernet with SFP optics
PIC 1	0 through 3	40-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 100-Gigabit Ethernet Default: Active ports operate in 100-Gigabit Ethernet mode.
PIC 2	38 and 39	200-Gigabit Ethernet



NOTE: For ACX5448M, prior to 20.4R1, when you configure a port speed, this reports a pic bounce alarm as a default behavior. From 20.4R1 onwards, there is no pic bounce alarm.

Table 16: Port Speed on ACX5448-M

PIC	Port	Port Speed Supported
PIC 0	0 through 43	10-Gigabit Ethernet with SFP+ optics 1-Gigabit Ethernet with SFP optics
PIC 1	0 through 3	100-Gigabit Ethernet with QSFP28 optics 4x25-Gigabit Ethernet (with breakout cables and channelization). 40-Gigabit Ethernet with QSFP+ optics 4x10-Gigabit Ethernet (with breakout cables and channelization).

Use the speed command to set the speed on tri-rate copper SFP port. For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Table 17: Interface Naming Convention for ACX5448

PIC	Interface type	ACX5448 Interfaces	ACX5448-D	ACX5448-M
PIC 0	1-Gigabit/10-Gigabit Ethernet interface (SFP+ or SFP ports)	xe-0/0/0 through xe-0/0/47	xe-0/0/0 through xe-0/0/35	xe-0/0/0 through xe-0/0/43

Table 17: Interface Naming Convention for ACX5448 (Continued)

PIC	Interface type	ACX5448 Interfaces	ACX5448-D	ACX5448-M
PIC 1	100-Gigabit Ethernet interface (QSFP28 ports)	et-0/1/48 :[0-3] through et-0/1/51 :[0-3]	et-0/1/36 :[0-3] and et-0/1/37 :[0-3]	et-0/1/44 :[0-3] through et-0/1/49 :[0-3]
PIC 2	200-Gigabit Ethernet interface (CFP2-DCO ports)	Not Supported	ot-0/2/38 and ot-0/2/39	Not Supported

For information on multiplexing on ACX5448-D routers, see [Multiplexing on ACX5448-D Routers](#).

Port Speed on ACX7020

ACX7020 has a fixed PIC with 20 ports that support multiple port speeds. All the interfaces of 1G, 10G and 25G speeds are named as et-x/y/z. The default port speed is 10G on all ports

For information about the line card, see ACX7020 Universal Metro Router Hardware Guide.

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs, see [Hardware Compatibility Tool \(HCT\)](#).

Table 18: Port Speed for ACX7020

PIC	Port Number	Port Speed Supported
PIC 0	0-3	1-Gigabit Ethernet 10-Gigabit Ethernet 25-Gigabit Ethernet
	4-19	1-Gigabit Ethernet 10-Gigabit Ethernet

Table 19: Interface Naming Convention for ACX7020

Optics Type	Interface Speed	Naming Convention
SFP	1-Gigabit Ethernet	et-x/y/z
SFP+	1x10-Gigabit Ethernet	et-x/y/z
SFP28	1x25-Gigabit Ethernet	et-x/y/z

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" [on page 177](#). For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" [on page 178](#).

Port Speed on ACX7024 and ACX7024X Router

On non-channelized interfaces, the prefix `et-` is used irrespective of the speed configured, default speed is assigned to the interface if you do not configure the speed. On channelized interfaces, the prefix `et-` is used irrespective of the speed configured. All channelized interfaces have the same speed. You cannot configure an individual speed for each channelized interface. You can configure the number of channelized interfaces by using the `number-of-sub-ports` command. When you change or remove, the speed or number-of-sub-ports of the port, the interfaces will be deleted and re-created for that port.

You can create Interfaces irrespective of the physical presence of optics. If the plugged in optics does not match the interface speed, the interfaces are marked down. You can configure port profiles in the command line interface without the physical presence of an FPC.

If an invalid port profile configuration is detected while booting a FPC, an alarm is generated. Also, the default port profile is selected for that PIC. Also, if the port profile configuration is changed while the FPC is up and running, and the new configuration is invalid, an alarm is generated. The existing port profile configured continues to be used for that PIC.

For information on ACX7024 and ACX7024X routers, see ACX7024 and ACX7024X Universal Metro Router Hardware Guide.

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#)

Table 20: Port speed for ACX7024 and ACX7024X routers

PIC	Port Number	Port Speed Supported
PIC 0	0-3	10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 100-Gigabit Ethernet
PIC 0	4-27	1-Gigabit Ethernet 10-Gigabit Ethernet 25-Gigabit Ethernet Default: 10-Gigabit Ethernet
	27	If you configure Precision Time Protocol (PTP), port 27 is disabled.

Table 21: Interface Naming Convention for ACX7024 and ACX7024X routers

Optics Type	Interface type	Interfaces
SFP	1-Gigabit Ethernet interface (SFP ports)	et-0/0/4 to et-0/0/27
SFP+	1 x 10-Gigabit Ethernet interface	et-0/0/4 to et-0/0/27
SFP28	1 x 25-Gigabit Ethernet interface	et-0/0/4 to et-0/0/27
QSFP28	1 x 100-Gigabit Ethernet interface	et-0/0/0 to et-0/0/3
	4 x 25-Gigabit Ethernet interface (Channelized Interface)	et-0/0/0:0 - et-0/0/3:0
		et-0/0/0:1 - et-0/0/3:1

Table 21: Interface Naming Convention for ACX7024 and ACX7024X routers (Continued)

Optics Type	Interface type	Interfaces
		et-0/0/0:2 - et-0/0/3:2
		et-0/0/0:3 - et-0/0/3:3
QSFP+	1 x 40-Gigabit Ethernet interface	et-0/0/0 to et-0/0/3
	4 x 10-Gigabit Ethernet interface (Channelized Interface)	et-0/0/0:0 - et-0/0/3:0
		et-0/0/0:1 - et-0/0/3:1
		et-0/0/0:2 - et-0/0/3:2
		et-0/0/0:3 - et-0/0/3:3

For information about how to configure the speed of a port for a non-channelized interface using the interfaces hierarchy, see ["Configure Port Speed for Non-Channelized Interfaces" on page 179](#). For information about how to configure the speed of a port for a channelized interface using the new interfaces hierarchy, see ["Configure Port Speed for Channelized Interfaces" on page 180](#).

Port Speed on ACX7100 Router

IN THIS SECTION

- [Port Speed on ACX7100-48L Router | 191](#)
- [Port Speed on ACX7100-32C Router | 194](#)

For information about ACX7100-48L router, see ACX7100-48L Universal Metro Router Hardware Guide.

For information on ACX7100-32C routers, see ACX7100-32C Universal Metro Router Hardware Guide.

Port Speed on ACX7100-48L Router

To view the supported transceivers, optical interfaces, and DAC cables on ACX7100-48L, see [Hardware Compatibility Tool \(HCT\)](#).

Use the speed command to set the speed on tri-rate SFP port. Use the show chassis pic fpc-slot 0 pic-slot 0 command to view supported port speeds. When optics does not support configured speed, an alarm is raised.

When you change the speed of a port within a port group, you'll observe a link flap on all ports of the port group.

Table 22: Port Speed for ACX7100-48L Router

PIC	Port Number	Supported Port Speed
PIC 0	port 0-46	10-Gigabit Ethernet
		25-Gigabit Ethernet
		50-Gigabit Ethernet
	port 47	10-Gigabit Ethernet
	If you configure Precision Time Protocol (PTP), port 47 is disabled.	25-Gigabit Ethernet 50-Gigabit Ethernet
	port 48-53	400-Gigabit Ethernet 1x40-Gigabit Ethernet 1x100-Gigabit Ethernet 4x100-Gigabit Ethernet 2x100-Gigabit Ethernet 1x100-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet

Table 23: ACX7100-48L Router Details and Description

Details	Description
Port speed configuration	Supports interfaces [edit interfaces interface-name] hierarchy.
MTU size	ACX7100-48L router WAN interfaces support maximum MTU size of 9996 bytes for transit traffic.
Forward Error Correction (FEC) support	<p>By default, FEC119 (KP4) is enabled on 400-Gigabit Ethernet and 200-Gigabit Ethernet interfaces.</p> <p>The FEC mode is based on the type of optics connected. When you configure the FEC mode, you must configure all interfaces on the port with the same FEC mode. For example, if you channelize a 40-Gigabit Ethernet port to operate as four 10-Gigabit Ethernet interfaces, you must configure each of the four 10-Gigabit Ethernet interfaces with either all FEC74 or all FEC91-KR4.</p> <p>For information about FEC support, see <i>fec (gigether)</i>.</p>

Table 24: Interface Naming Convention for ACX7100-48L Router

PIC	Interface Type	Interfaces
PIC 0	10-Gigabit Ethernet interface 25-Gigabit Ethernet interface 50-Gigabit Ethernet interface	et-0/0/0- et-0/0/47

Table 24: Interface Naming Convention for ACX7100-48L Router *(Continued)*

PIC	Interface Type	Interfaces
	400-Gigabit Ethernet	et-0/0/48-et-0/0/53
	1x40-Gigabit Ethernet	
	1x100-Gigabit Ethernet	
	4x100-Gigabit Ethernet	
	2x100-Gigabit Ethernet	
	1x100-Gigabit Ethernet	
	8x50-Gigabit Ethernet	
	4x10-Gigabit Ethernet	
	4x25-Gigabit Ethernet	
	interface (6 QSFP56-DD ports)	

ACX7100-48L has twelve Port Macro (PMs). Six of the twelve PMs are for port 0-47. Each PM has eight ports. By default, port 0-47 support 10G.

Port Macro does not support the combination of port speed, 25Gbps and 50Gbps. For example, for ports 0-7, you can specify the speed for port 4 as 50G and for port 2 as 10G. Alternatively, you can specify the speed of port 4 as 25G and for port 2 as 10G. When you configure the speed for port 4 as 50G and for port 2 as 25G and commit the configuration, an alarm is generated.

Table 25: ACX7100-48L Port Macro Port Speed Configuration Router

Per Port Macro Port	Port Speed Configuration			Supported
	10-Gigabit Ethernet	25-Gigabit Ethernet	50-Gigabit Ethernet	
Ports (0-7)	Yes	Yes	Yes	Yes
(8-15)				
(16-23)				
(24-31)				
(32-39)				
(40-46)				
47	Yes	Yes	No	Yes

Port Speed on ACX7100-32C Router

To view the supported transceivers, optical interfaces, and DAC cables on ACX7100-32C, see [Hardware Compatibility Tool \(HCT\)](#).

If an even port is configured in channelized mode (4x10G or 4x25G), then you cannot use the subsequent odd port and need to configure the port as unused.

When you change the speed of a port within a port group, you'll observe a link flap on all ports of the port group.

[

Table 26: Port Speed for ACX7100-32C

PIC	Port Number	Port Speed Supported
PIC 0	Port 0-31	100-Gigabit Ethernet
		40-Gigabit Ethernet
		25-Gigabit Ethernet
	Port 31	10-Gigabit Ethernet
		1 Gigabit Ethernet
		When you configure PTP, port 31 can be used to connect grand master (GM) clock with 1G speed.
	Port 32-35	400-Gigabit Ethernet
		1x40-Gigabit Ethernet
		1x100-Gigabit Ethernet
		4x100-Gigabit Ethernet
		2x100-Gigabit Ethernet
		1x100-Gigabit Ethernet
		4x10-Gigabit Ethernet
		4x25-Gigabit Ethernet

Table 27: Port Group for ACX7100-32C

PIC	Port Number	Supported Port Speed
PIC 0	(0-3), (4-7), (8-11), (12-15), (16-19), (20-23), (24-27), and (28-31)	1x10-Gigabit Ethernet 4x10-Gigabit Ethernet 1x25-Gigabit Ethernet 4x25-Gigabit Ethernet 1x40-Gigabit Ethernet 2x50-Gigabit Ethernet 2x100-Gigabit Ethernet

Table 28: Interface Naming Convention for ACX7100-32C

PIC	Interface type	Interfaces
PIC 0	Port numbers: 0/ 2/ 4/ 6/ 8/ 10/ 12/ 14/ 16/ 18/ 20/ 22/ 24/ 26/ 28/ 30: 1x40-Gigabit/1x100-Gigabit /4x10-Gigabit/1x25-Gigabit/4x25-Gigabit/1x10-Gigabit Ethernet interface Port numbers: 1/ 3/ 5/ 7/ 9/ 11/ 13/ 15/ 17/ 19/ 21/ 23/ 25/ 27/ 29: 1x100-Gigabit/1x40-Gigabit/1x25-Gigabit/1x10-Gigabit Ethernet interface (32 QSFP28 ports)	et-0/0/0- et-0/0/31
	400-Gigabit/1x40-Gigabit/1x100-Gigabit/4x100-Gigabit/2x100-Gigabit/1x100-Gigabit/8x50-Gigabit/4x10-Gigabit/4x25-Gigabit Ethernet interface (4 QSFP56-DD ports)	et-0/0/32- et-0/0/35

For information about how to configure the speed of a port for a non-channelized interface using the interfaces hierarchy, see ["Configure Port Speed for Non-Channelized Interfaces" on page 179](#). For

information about how to configure the speed of a port for a channelized interface using the new interfaces hierarchy, see ["Configure Port Speed for Channelized Interfaces" on page 180](#).

Port Speed on ACX7332 and ACX7348 Routers

The ACX7332 router has one fixed FPC, 32 SFP, and 8 QSFP ports that support 1- Gbps, 10-Gbps, 25-Gbps, 40-Gbps, and 100-Gbps speeds. Active ports are mentioned in the below table.

The ACX7348 router has one fixed FPC, 48 SFP, and 8 QSFP ports that support 1- Gbps, 10-Gbps, 25-Gbps, 40-Gbps, and 100-Gbps speeds. Active ports are mentioned in the below table.

SFP port numbers: et-0/0/0 to et-0/0/23 and et-0/0/32 to et-0/0/39 supports 1Gbps, 10Gbps, and 25Gbps. QSFP port numbers: et-0/0/24 to et-0/0/31 supports 40Gbps, 4x10Gbps, 4x25Gbps and 100Gbps. FEC 91 supports 100Gbps, and FEC 74 on 25Gbps DAC.

For information on ACX7332 router, see ACX7332 Universal Metro Router Hardware Guide.

For information on ACX7348 router, see ACX7348 Universal Metro Router Hardware Guide.

For information on the line card, see [Hardware Compatibility Tool \(HCT\)](#).

Table 29: Port Speed for ACX7332 Fixed FPC Interface

Slot	Port Number	Supported Port Speed
0	0 - 23	1x1-Gigabit Ethernet
		1x10-Gigabit Ethernet
		1x25-Gigabit Ethernet
0	24 - 31	1x40-Gigabit Ethernet
		1x100-Gigabit Ethernet
		4x10-Gigabit Ethernet
		4x25-Gigabit Ethernet
0	32 - 39	1x1-Gigabit Ethernet
		1x10-Gigabit Ethernet
		1x25-Gigabit Ethernet

Table 30: Port Speed for ACX7348 Fixed FPC Interface

Slot	Port Number	Supported Port Speed
0	0 - 23	1x1-Gigabit Ethernet
		1x10-Gigabit Ethernet
		1x25-Gigabit Ethernet
0	24 - 31	1x40-Gigabit Ethernet
		1x100-Gigabit Ethernet
		4x10-Gigabit Ethernet
		4x25-Gigabit Ethernet
0	32 - 55	1x1-Gigabit Ethernet
		1x10-Gigabit Ethernet
		1x25-Gigabit Ethernet

The 6xQSFPDD FPC Line Card (ACX7332 and ACX7348) has two QSFPDD ports (Port 0 and 2) and four QSFP ports (Port 1, 3, 4 and 5). Active ports are mentioned in the below table.

Table 31: Port Speed for 6xQSFPDD FPC Line Card (ACX7332 and ACX7348)

Slot	Port Number	Supported Port Speed
1,2	0, 2	4x10-Gigabit Ethernet
		1x40-Gigabit Ethernet
		4x25-Gigabit Ethernet
		1x100-Gigabit Ethernet
		2x100-Gigabit Ethernet
		4x100-Gigabit Ethernet
		1x400-Gigabit Ethernet

Table 31: Port Speed for 6xQSFPDD FPC Line Card (ACX7332 and ACX7348) (Continued)

Slot	Port Number	Supported Port Speed
1,2	1,3,4,5	4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 1x40-Gigabit Ethernet 1x100-Gigabit Ethernet
3	0,1,2	4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 1x40-Gigabit Ethernet 1x100-Gigabit Ethernet

In a 6xQSFPDD FPC line card, all the ports supports different speed capability. The table shows the supported speed combinations.

Table 32: Port Group for 6xQSFPDD FPC Line Card (ACX7332 and ACX7348)

Slot	Port Group	Supported Port Speed Combination		
		Port 0 or Port 2	Port 1 or Port 4	Port 3 or Port 5
1, 2	(0,1,3) (2,4,5)	1x400G 4x100G	Unused	Unused
1,2	(0,1,3) (2,4,5)	8x25G	Unused	Unused
1,2	(0,1,3) (2,4,5)	2x100G 4x25G 4x10G 1x100G 1x40G	1x100G 1x40G	1x100G 1x40G

Table 32: Port Group for 6xQSFPDD FPC Line Card (ACX7332 and ACX7348) (Continued)

Slot	Port Group	Supported Port Speed Combination		
		Port 0 or Port 2	Port 1 or Port 4	Port 3 or Port 5
1,2	(0,1,3) (2,4,5)	2x100G	4x25G	Unused
		4x25G	4x10G	
		4x10G		
		1x100G		
		1x40G		
1,2	(0,1,3) (2,4,5)	2x100G	Unused	4x25G
		4x25G		4x10G
		4x10G		4x50G
		1x100G		
		1x40G		
3	(0,1,2)	1x100G	1x100G	Unused
		4x25G	4x25G	
		1x40G	1x40G	
		4x10G	4x10G	

The 16xSFP56 FPC line card (ACX7332 and ACX7348) has 16 SFP56 ports (Port 0 to 15) that support 1- Gbps, 10-Gbps, 25-Gbps, 40-Gbps, and 50-Gbps speeds. Active ports are mentioned in the below table.

Table 33: Port Speed for 16xSFP56 FPC Line Card (ACX7332 and ACX7348)

Slot	Port Number	Supported Port Speed
1,2	0-7	1x10-Gigabit Ethernet 1x25-Gigabit Ethernet OR 1x10-Gigabit Ethernet 1x50-Gigabit Ethernet NOTE: You cannot combine 1x25-Gigabit Ethernet and 1x50-Gigabit Ethernet.
1,2	8-15	1x10-Gigabit Ethernet 1x25-Gigabit Ethernet OR 1x10-Gigabit Ethernet 1x50-Gigabit Ethernet NOTE: You cannot combine 1x25-Gigabit Ethernet and 1x50-Gigabit Ethernet.
3	0-13	4x10-Gigabit Ethernet 1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet

For information about how to configure the speed of a port for a non-channelized interface using the interfaces hierarchy, see ["Configure Port Speed for Non-Channelized Interfaces" on page 179](#). For information about how to configure the speed of a port for a channelized interface using the new interfaces hierarchy, see ["Configure Port Speed for Channelized Interfaces" on page 180](#).

Port Speed on ACX7509 Router

IN THIS SECTION

- [Port Speed for ACX7509-FPC-20Y Line Card | 202](#)
- [Port Speed on ACX7509-FPC-16C Line Card | 204](#)
- [Port Speed on ACX7509-FPC-4CD Line Card | 208](#)

For information on ACX7509 router, see ACX7509 Universal Metro Router Hardware Guide.

For information on the line card, see [Hardware Compatibility Tool \(HCT\)](#).

Port Speed for ACX7509-FPC-20Y Line Card

The ACX7509-FPC-20Y line card has 20 SFP56 ports that support 1-Gbps, 10-Gbps, and 25-Gbps speeds.

When you attempt to configure an invalid or unsupported speed for any port on the router, the port continues to operate with the default speed. An alarm is raised indicating that the port speed configuration is invalid.

Table 34: Port Speed for ACX7509-FPC-20Y Line Card

Slot	Port Number	Supported Port Speed
0, 2, 3, 4, 6, and 7	Ports 0 through 19	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet Default Speed: 25-Gigabit Ethernet
1 and 5	Ports 0 through 19	1x10-Gigabit Ethernet 1x25-Gigabit Ethernet Default Speed: 25-Gigabit Ethernet

Ports within a port group can support mixed speeds. For instance, on slot 0, within port group (0-3), each port can support a different port speed (1-Gbps, 10-Gbps, or 25-Gbps). When you change the speed of a port within a port group, you'll observe a link flap on all ports of the port group.

Table 35: Port Group for ACX7509-FPC-20Y Line Card

Slot	Port Number	Supported Port Speed
0, 2, 3, 4, 6, 7	(0-3), (4-7), (8-11), (12-15), and (16-19)	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet
1 and 5	(0-7), (8-16), and (16-19)	1x25-Gigabit Ethernet 1x10-Gigabit Ethernet

Table 36: ACX7509-FPC-20Y Line Card Details and Description

Details	Description
Autonegotiation	Autonegotiation is supported only on the following channelization: <ul style="list-style-type: none">Twenty 1-Gigabit Ethernet interface in ACX7509-FPC-20Y
Channelization	Not supported
Port speed configuration	Supports interfaces [edit interfaces interface-name] hierarchy.
MTU size	ACX7509-FPC-20Y WAN interfaces support maximum MTU size of 9996 bytes for transit traffic.
Forward Error Correction (FEC) support	FEC91 is enabled on 25-Gigabit optics interfaces. FEC74 is enabled on 25-Gigabit DAC cables. For information about FEC support, see <i>fec (gigether)</i> .

Table 37: Interface Naming Convention for ACX7509-FPC-20Y Line Card

PIC	Interface Type	Interfaces
PIC 0	1-Gigabit Ethernet	et-0/0/0-0/0/19
	10-Gigabit Ethernet	
	25-Gigabit Ethernet	et-0/0/0-0/0/19

Port Speed on ACX7509-FPC-16C Line Card

The ACX7509-FPC-16C line card has 16 QSFP28 ports that support 100-Gbps and 40-Gbps speeds. You can configure the speed of active ports only. You can configure 0,1,4,5,8,9, and 12,13 ports using breakout cables with four 25-Gbps and four 10-Gbps speed.

If you channelize (4x10G or 4x25G) and non-channelize speed on adjacent ports, an alarm is generated indicating that the port speed configuration is invalid. Ensure you channelize adjacent ports and configure the unintended next port as unused. For example, to channelize Port 0, you must channelize both Port 0 and Port 1 or you must configure Port 1 as Unused.

When you attempt to configure an unused port, an alarm is generated indicating that the slot is unsupported. For instance, if you attempt to configure the speed on Port 2, an alarm is generated indicating that the slot is unsupported.

Table 38: Port Speed for ACX7509-FPC-16C Line Card

Slot	Port Number	Supported Port Speed
0	Port 0,1,4,5,8,9,12, and 13	1x40-Gigabit Ethernet 4x10-Gigabit Ethernet (using breakout cables only) 4x25-Gigabit Ethernet (using breakout cables only) Default Speed: 1x100-Gigabit Ethernet
1	Ports 0 through 15	Default Speed: 1x100-Gigabit Ethernet

Table 38: Port Speed for ACX7509-FPC-16C Line Card *(Continued)*

Slot	Port Number	Supported Port Speed
2	Port 0,1,4,5,8,9,12, and 13	1x40-Gigabit Ethernet 4x10-Gigabit Ethernet (using breakout cables only) 4x25-Gigabit Ethernet (using breakout cables only) Default Speed: 1x100-Gigabit Ethernet
3	Port 0,1,4,5,8,9,12, and 13	1x40-Gigabit Ethernet 4x10-Gigabit Ethernet (using breakout cables only) 4x25-Gigabit Ethernet (using breakout cables only) Default Speed: 1x100-Gigabit Ethernet
4	Port 0,1,4,5,8,9,12, and 13	1x40-Gigabit Ethernet 4x10-Gigabit Ethernet (using breakout cables only) 4x25-Gigabit Ethernet (using breakout cables only) Default Speed: 1x100-Gigabit Ethernet
5	Ports 0 through 15	Default Speed: 1x100-Gigabit Ethernet
6	Port 0,1,4,5,8,9,12, and 13	1x40-Gigabit Ethernet 4x10-Gigabit Ethernet (using breakout cables only) 4x25-Gigabit Ethernet (using breakout cables only) Default Speed: 1x100-Gigabit Ethernet

Table 38: Port Speed for ACX7509-FPC-16C Line Card *(Continued)*

Slot	Port Number	Supported Port Speed
7	Port 0,1,4,5,8,9, and 12	1x40-Gigabit Ethernet 4x10-Gigabit Ethernet (using breakout cables only) 4x25-Gigabit Ethernet (using breakout cables only) Default Speed: 1x100-Gigabit Ethernet
slots (0,2,3,4,6, and 7)	Ports: 2,3,6,7,10,11,14, and 15	Not supported

Table 39: ACX7509-FPC-16C Line Card Details and Description

Details	Description
Channelization	You can channelize the ports operating at: <ul style="list-style-type: none">• 100-Gbps speed to 4x25-Gbps speed• 40-Gbps speed to 4x10-Gbps speed
Port speed configuration	Supports interfaces [edit interfaces interface-name] hierarchy.
MTU size	ACX7509-FPC-16C WAN interfaces support maximum MTU size of 9996 bytes for transit traffic.

Table 39: ACX7509-FPC-16C Line Card Details and Description (Continued)

Details	Description
Forward Error Correction (FEC) support	<p>FEC91 is enabled on 100-Gigabit and 25-Gigabit Ethernet interfaces. FEC91 is enabled by default on 25-Gigabit Ethernet interfaces.</p> <p>On the Direct Attach Copper (DAC) connection of the 25-Gigabit Ethernet interface, FEC74 is enabled.</p> <p>The FEC mode is based on the type of optics connected. When you configure the FEC mode, you must configure all interfaces on the port with the same FEC mode. For example, if you channelize a 100-Gigabit Ethernet port to operate as four 25-Gigabit Ethernet interfaces, you must configure each of the four 25-Gigabit Ethernet interfaces with either all FEC74 or all FEC91-KR4.</p> <p>For information about FEC support, see <i>fec (gigether)</i>.</p>

Table 40: Interface Naming Convention for ACX7509-FPC-16C Line Card

Slot	1x40-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface 4x25-Gigabit Ethernet Interface
Slot 1	et-0/0/0	et-0/0/0	et-0/0/0:[0-3]
	et-0/0/1	et-0/0/1	et-0/0/1:[0-3]
	et-0/0/2	et-0/0/2	et-0/0/2
	et-0/0/3	et-0/0/3	et-0/0/3
	et-0/0/4	et-0/0/4	et-0/0/4:[0-3]
	et-0/0/5	et-0/0/5	et-0/0/5:[0-3]
	et-0/0/6	et-0/0/6	et-0/0/6

Table 40: Interface Naming Convention for ACX7509-FPC-16C Line Card (*Continued*)

Slot	1x40-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface 4x25-Gigabit Ethernet Interface
	et-0/0/7	et-0/0/7	et-0/0/7
	et-0/0/8	et-0/0/8	et-0/0/8:[0-3]
	et-0/0/9	et-0/0/9	et-0/0/9:[0-3]
	et-0/0/10	et-0/0/10	et-0/0/10
	et-0/0/11	et-0/0/11	et-0/0/11
	et-0/0/12	et-0/0/12	et-0/0/12:[0-3]
	et-0/0/13	et-0/0/13	et-0/0/13:[0-3]
	et-0/0/14	et-0/0/14	et-0/0/14
	et-0/0/15	et-0/0/15	et-0/0/15

Port Speed on ACX7509-FPC-4CD Line Card

The ACX7509-FPC-4CD line card has 4 QSFP56-DD ports that support 4x100-Gbps and 400-Gbps speeds. ACX7509-FPC-4CD line card is supported only in FPC slot 1 and 5.

Table 41: Port Speed for ACX7509-FPC-4CD Line Card

Slot	Port Number	Supported Port Speed
1 and 5	Ports: 0 through 3	1x400-Gigabit Ethernet 4x100-Gigabit Ethernet (using breakout cables only) Default Speed: 400-Gigabit Ethernet
0, 2, 3, 4, 6, 7	Not supported	Not supported

Table 42: ACX7509-FPC-4CD Line Card Details and Description

Details	Description
Channelization	You can channelize the ports operating at: <ul style="list-style-type: none">• 400-Gbps speed to 4x100-Gbps speed.
Port speed configuration	Supports interfaces [edit interfaces interface-name] hierarchy.
MTU size	ACX7509-FPC-4CD WAN interfaces support maximum MTU size of 9996 bytes for transit traffic.
Forward Error Correction (FEC) support	FEC119 is enabled on 400-Gigabit Ethernet interfaces. FEC91 is enabled on 100-Gigabit Ethernet interfaces. For information about FEC support, see <i>fec (gigether)</i> .

Table 43: Interface Naming Convention for ACX7509-FPC-4CD Line Card

PIC	Interface Type	Interfaces
PIC 0	1x400-Gigabit Ethernet	et-0/0/0- et-0/0/3
	4x100-Gigabit Ethernet	et-0/0/0:[0-3]- et-0/0/3:[0-3]

For information about how to configure the speed of a port for a non-channelized interface using the interfaces hierarchy, see ["Configure Port Speed for Non-Channelized Interfaces" on page 179](#). For information about how to configure the speed of a port for a channelized interface using the new interfaces hierarchy, see ["Configure Port Speed for Channelized Interfaces" on page 180](#).

RELATED DOCUMENTATION

*request security system-keystore unlink
speed (Ethernet)*

Port Speed on MX Routers

IN THIS SECTION

- [Port Speed on MX204 | 211](#)
- [Port speed on MX301 Router | 216](#)
- [Port speed on MX304 Router | 221](#)
- [Port Speed on MPC7E-MRATE | 224](#)
- [Port Speed on MPC7E-10G | 227](#)
- [Port Speed on MIC-MRATE | 229](#)
- [Port Speed on MIC-3D-10GE-SFP-E | 233](#)
- [Port Speed on MX10003 MPC | 234](#)
- [Port Speed on MX10K-LC4800 | 240](#)
- [Port Speed on MIC-MACSEC-20GE | 244](#)
- [Port Speed on MPC10E-10C-MRATE | 245](#)
- [Port Speed on MPC10E-15C-MRATE | 250](#)
- [Port Speed on MX2K-MPC11E | 256](#)
- [Port Speed on MX10K-LC480 | 266](#)
- [Port Speed on MX10K-LC2101 | 268](#)
- [Port Speed on MX10K-LC9600 | 272](#)

In this topic we have information related to the port speed on a MX routers and line cards, support for multiple port speed details, guidelines and how to configure the port speed.

Port Speed on MX204

On an MX204 router, there are two line card— MIC (PIC1) and fixed-port PIC (PIC0). You must configure both the PICs and all the associated ports, under the [edit chassis] hierarchy. Configuring ports on only one of the PICs results in an invalid configuration.

You cannot configure port speed at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations. By default, the MX204 router comes up with PIC mode where all the interfaces operate at the same speed of 10-Gbps, and both the PICs (PIC0 and PIC1) operate at 10-Gbps speed. To use different port speeds, you must first switch to the port mode. When you switch modes, either from PIC mode to port mode or from port mode to PIC mode, the PIC restarts automatically.

For information on the router, see [MX204 Universal Routing Platform Hardware Guide](#).

To view the supported transceivers, optical interfaces, and DAC cables on MX204, see [Hardware Compatibility Tool \(HCT\)](#).

To change the default speed, you must select a port and configure a different port speed on it, and reset both the PICs. For example, select 40GE or 100GE on PIC0 and 10GE on PIC1. For this configuration to take effect, you must reset both PICs.

Table 44: Port Speed for MX204

PIC	Port	Port Speed Supported
PIC 0	0-3	4x10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 4x10-Gigabit Ethernet Supports 1-Gbps speed on 10 Gigabit Ethernet ports (Junos OS Release 18.1 onwards)

Table 44: Port Speed for MX204 (*Continued*)

PIC	Port	Port Speed Supported
PIC 1	0-7	10-Gigabit Etherner Supports 1-Gbps speed on 10 Gigabit Ethernet ports (Junos OS Release 18.1 onwards)

When you choose an existing port speed configuration with an invalid port speed configuration, an alarm is generated indicating that the port speed configuration is invalid. You cannot configure 1-Gbps speed at PIC level and port level. But, you can configure the port configured at 10-Gbps speed to operate at 1-Gbps speed by using the speed statement. After you commit the configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed, but the show interface command displays the speed configuration (operating port speed) as 1GE.

Table 45: Maximum number of ports Configurable at PIC or Port level (MX204)

Type of Ports	PIC Mode (maximum configurable ports on PIC0 and PIC1)	Port Mode (maximum configurable ports on PIC0 and PIC1)
10/1 Gigabit Ethernet Ports	24 16 ports from PIC0 and 8 ports from PIC1	20 12 ports from PIC0 and 8 ports from PIC1
40 Gigabit Ethernet Ports	4 4 ports from PIC0	4
100 Gigabit Ethernet Ports	4 4 ports from PIC0	4

When you configure the port speed at the port level, only the configured ports are active while other ports are disabled. If you change an existing port speed configuration at the port level, reset the PIC for the configuration to take effect.

Table 46: Valid Port Speed Combinations at Port Level (MX204)

Port Mode												
PIC0				PIC1								
100	100	100	100	0	Configure the number of active ports as 0							
100	100	100	10/40	0	Configure the number of active ports as 0							
100	100	10/40	10/40	10	10	10	10	10	10	10	10	10
100	10/40	10/40	10/40	10	10	10	10	10	10	10	10	10
10/40	10/40	10/40	10/40	0	Configure the number of active ports as 0							

In PIC mode, the MX204 router does not support heterogeneous mode. That is, in PIC mode if 40-Gbps or 100-Gbps speed is configured on PIC 0, then the number-of-ports on PIC 1 must be configured to 0 only. The heterogeneous mode is supported only on port mode. When you change an existing port speed configuration at the PIC level, the MPC automatically resets the PIC.

Table 47: Valid Port Speed Combinations at PIC level (MX204)

PIC Mode				
PIC0 (with four multirate ports)				PIC1 (with eight 10-Gigabit Ethernet ports)
100	100	100	100	0 Configure the number of active ports as 0

Table 47: Valid Port Speed Combinations at PIC level (MX204) (Continued)

PIC Mode				
40	40	40	40	0 Configure the number of active ports as 0
10	10	10	10	0 Configure the number of active

Table 48: Interface Naming Conventions (MX204)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe-x/0/0:[0-3]	et-x/0/0:[0-3]	et-x/0/0:[0-3]
	xe-x/0/1:[0-3]	et-x/0/1:[0-3]	et-x/0/1:[0-3]
	xe-x/0/2:[0-3]	et-x/0/2:[0-3]	et-x/0/2:[0-3]
	xe-x/0/3:[0-3]	et-x/0/3:[0-3]	et-x/0/3:[0-3]
1	xe-x/1/0:[0-3]	et-x/1/0:[0-3]	et-x/1/0:[0-3]
	xe-x/1/1:[0-3]	xe-x/1/1:[0-3]	xe-x/1/1:[0-3]
	xe-x/1/2:[0-3]	xe-x/1/2:[0-3]	xe-x/1/2:[0-3]
	xe-x/1/3:[0-3]	xe-x/1/3:[0-3]	xe-x/1/3:[0-3]
	xe-x/1/4:[0-3]	xe-x/1/4:[0-3]	xe-x/1/4:[0-3]

Table 48: Interface Naming Conventions (MX204) (Continued)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
	xe-x/1/5:[0-3]	xe-x/1/5:[0-3]	xe-x/1/5:[0-3]
	xe-x/1/6:[0-3]	xe-x/1/6:[0-3]	xe-x/1/6:[0-3]
	xe-x/1/7:[0-3]	xe-x/1/7:[0-3]	xe-x/1/7:[0-3]

Table 49: Active Physical Ports (MX204) at PIC Level

PIC	Number of Ports Configured	Active Ports		
		10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
PIC0		-	-	-
	0	0	0	0
	1	0, 1	0, 1	0, 1
	2	0, 1, 2	0, 1, 2	0, 1, 2
PIC1	0	-	-	-
	1	0	-	-
	2	0, 1	-	-
	3	0, 1, 2	-	-

Table 49: Active Physical Ports (MX204) at PIC Level (*Continued*)

PIC	Number of Ports Configured	Active Ports		
	4	0, 1, 2, 3	-	-
	5	0, 1, 2, 3, 4	-	-
	6	0, 1, 2, 3, 4, 5	-	-
	7	0, 1, 2, 3, 4, 5, 6	-	-

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port speed on MX301 Router

MX301 use port profiles to manage the PIC ports. A port profile selects a set of ports that will be active in a PIC and the port speed. The port profiles can be configured at a PIC level as well as at a per-port level. The port profile configuration allows you to configure the port speed at PIC level. All the ports supporting that port speed will be active by default.

The number-of-ports CLI configuration command is not supported.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs, see [Hardware Compatibility Tool \(HCT\)](#).

For information on MX301, see Protocols and Application supported by the MX301.

Only the ports specified in the CLI configuration will be treated as active ports and the interfaces are created for active ports only. When a port profile configuration is changed, the interfaces corresponding to the affected ports will be deleted and re-created. No PIC bounce or reset of MPC will happen for the port profile configuration changes.

Port speeds 10G and 25G need number-of-subport which is either 1 or 4. For 1G, 50G, and 100G it is single channel by default. For 10G and 25G we need to give number-of-subports as 1 or 4 or 8 as required.

Table 50: Port Speed for MX301

PIC	Port Number	Port Speed Supported
PIC 0 through PIC 2	0, 22	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 8x10-Gigabit Ethernet 1x25-Gigabit Ethernet 8x25-Gigabit Ethernet 8x50-Gigabit Ethernet 1x100-Gigabit Ethernet 2x100-Gigabit Ethernet 3x100-Gigabit Ethernet 4x100-Gigabit Ethernet 400-Gigabit Ethernet
	1,23	1x1-Gigabit Ethernet 4x1-Gigabit Ethernet 1x10-Gigabit Ethernet 4x10-Gigabit Ethernet 1x25-Gigabit Ethernet 4x25-Gigabit Ethernet 100-Gigabit Ethernet
	2, 24	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet 100-Gigabit Ethernet

Table 50: Port Speed for MX301 *(Continued)*

PIC	Port Number	Port Speed Supported
	3, 25	1x1-Gigabit Ethernet 4x1-Gigabit Ethernet 1x10-Gigabit Ethernet 4x10-Gigabit Ethernet 1x25-Gigabit Ethernet 4x25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet
	4, 6, 8, 10, 14, 16, 18, 20	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet 1x50-Gigabit Ethernet 100-Gigabit Ethernet
	5, 7, 9, 11, 15, 17, 19, 21	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet 1x50-Gigabit Ethernet

Table 50: Port Speed for MX301 (*Continued*)

PIC	Port Number	Port Speed Supported
	12, 13	4x1-Gigabit Ethernet 4x10-Gigabit Ethernet 8x10-Gigabit Ethernet 4x25-Gigabit Ethernet 8x25-Gigabit Ethernet 8x50-Gigabit Ethernet 100-Gigabit Ethernet 2x100-Gigabit Ethernet 3x100-Gigabit Ethernet 4x100-Gigabit Ethernet 400-Gigabit Ethernet

Table 51: Interface Naming Convention for MX301

PIC	Interface Type	Interfaces
PIC 0	1x1-Gigabit Ethernet Interface	ge-x/y/z:0
	4x1-Gigabit Ethernet Interface	ge-x/y/z:0 ge-x/y/z:1 ge-x/y/z:2 ge-x/y/z:3
	1x10-Gigabit Ethernet Interface	xe-x/y/z:0

Table 51: Interface Naming Convention for MX301 *(Continued)*

PIC	Interface Type	Interfaces
	4x10-Gigabit Ethernet Interface	xe-x/y/z:0 xe-x/y/z:1 xe-x/y/z:2 xe-x/y/z:3
	1x25-Gigabit Ethernet Interface	et-x/y/z:0
	4x25-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1 et-x/y/z:2 et-x/y/z:3
	1x40-Gigabit Ethernet Interface	et-x/y/z
	1x50-Gigabit Ethernet Interface	et-x/y/z
	8x50-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1 et-x/y/z:2 et-x/y/z:3 et-x/y/z:4 et-x/y/z:5 et-x/y/z:6 et-x/y/z:7
	1x100-Gigabit Ethernet Interface	et-x/y/z

Table 51: Interface Naming Convention for MX301 (Continued)

PIC	Interface Type	Interfaces
	2x100-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1
	4x100-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1 et-x/y/z:2 et-x/y/z:3
	3x100-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1 et-x/y/z:2
	2x200-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1
	1x400-Gigabit Ethernet Interface	et-x/y/z

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port speed on MX304 Router

Starting in Junos OS Release 23.1R1, the MX304 supports 1G speeds in addition to the previously supported speeds of 400G, 100G, 40G, 25G, and 10G. Also, supports 1G Class-of-Service (Cos) on 1G ports. By default, the MX304 router comes up with PIC mode where all the interfaces operate at the same speed of 100 Gbps. To use different port speeds, you must first switch to the port mode. When you switch modes, either from PIC mode to port mode or from port mode to PIC mode, the PIC restarts automatically. When a single 10 Gbps interface is selected, define the number-of-sub-ports value as 1. When you select 1 Gbps interface, do not configure number-of-sub-ports.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs, see [Hardware Compatibility Tool \(HCT\)](#).

For information on MX304, see [Protocols and Application supported by the MX304](#).

When you configure and commit the port speed simultaneously at the PIC and port levels, error messages are displayed. When you change an existing port speed configuration with an invalid port speed, an alarm is generated. The fpc continues to operate with the last available configuration. If the fpc is rebooted with an invalid configuration, the fpc move to the default mode.

When you change the port configuration at the port level, the interfaces corresponding to the affected port are deleted and then re-created.

Table 52: Port Speed for MX304

PIC	Port Number	Port Speed Supported
PIC 0 through PIC 2	2, 4, 10, 12	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet 1x100-Gigabit Ethernet Supports 1-Gbps speed (Junos OS Release 23.1R1 onwards) Default: All active ports operate in 100-Gigabit Ethernet mode.

Table 52: Port Speed for MX304 (*Continued*)

PIC	Port Number	Port Speed Supported
	0, 6, 8, 14	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet 1x100-Gigabit Ethernet 1x400-Gigabit Ethernet 4x100-Gigabit Ethernet 3x100-Gigabit Ethernet 2x100-Gigabit Ethernet Supports 1-Gbps speed (Junos OS Release 23.1R1 onwards)
	1, 3, 5, 7, 9, 11, 13, 15	1x1-Gigabit Ethernet 1x10-Gigabit Ethernet 1x25-Gigabit Ethernet 1x100-Gigabit Ethernet 4x1-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 1x40-Gigabit Ethernet Supports 1-Gbps speed (Junos OS Release 23.1R1 onwards)

Table 53: Interface Naming Convention for MX304

PIC	Interface Type	Interfaces
PIC 0	1x1-Gigabit Ethernet Interface	ge-0/0/0

Table 53: Interface Naming Convention for MX304 (Continued)

PIC	Interface Type	Interfaces
	1x10-Gigabit Ethernet Interface	xe-0/0/0:0
	4x10-Gigabit Ethernet Interface	xe-0/0/0:0 – xe-0/0/0:3
	1x25-Gigabit Ethernet Interface	et-0/0/0:0
	1x40-Gigabit Ethernet Interface	et-0/0/0
	1x100-Gigabit Ethernet Interface	
	1x400-Gigabit Ethernet Interface	
	4x25-Gigabit Ethernet Interface	et-0/0/0:0 – et-0/0/3:0
	4x100-Gigabit Ethernet Interface	
	8x50-Gigabit Ethernet Interface	et-0/0/0:0 – et-0/0/0:7
	2x100-Gigabit Ethernet Interface	et-0/0/0:0 – et-0/0/0:1

For information on how to configure the speed at the PIC level, see ["Configure Port Speed - PIC Level" on page 177](#). For information on how to configure the speed at the port level, see ["Configure Port Speed - Port Level" on page 178](#).

Port Speed on MPC7E-MRATE

By default, all ports in the MPC7E-10G MPC operate as 10 Gigabit Ethernet interfaces. When you configure the port speed, the PIC and MIC operate at the configured speed. An alarm is generated if you configure an invalid speed and when you change an existing port speed.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MPC7E-MRATE, see [Hardware Compatibility Tool \(HCT\)](#).

Table 54: Port Speed for MPC7E-MRATE

PIC	Port	Port Speed Supported
PIC 0 (Fixed-port PIC)	0, 1, 3, 4	40-Gigabit Ethernet
	2 and 5	100-Gigabit Ethernet
		Default: 10-Gigabit Ethernet
PIC 1 (Fixed-port PIC)	0, 1, 3,4	40-Gigabit Ethernet
	2 and 5	100-Gigabit Ethernet
		Default: 10-Gigabit Ethernet

[Table 55 on page 225](#) lists the naming conventions used for interfaces on MPC7E-MRATE for MX240, MX480, MX960, MX2010, and MX2020 routers.

Table 55: Interface Naming Convention for MPC7E-MRATE

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	<code>xe-x/0/0:[0-3]</code>	<code>et-x/0/0</code>	-
	<code>xe-x/0/1:[0-3]</code>	<code>et-x/0/1</code>	-
	<code>xe-x/0/2:[0-3]</code>	<code>et-x/0/2</code>	<code>et-x/0/2</code>
	<code>xe-x/0/3:[0-3]</code>	<code>et-x/0/3</code>	-
1	<code>xe-x/0/4:[0-3]</code>	<code>et-x/0/4</code>	-
	<code>xe-x/0/5:[0-3]</code>	<code>et-x/0/5</code>	<code>et-x/0/5</code>
	<code>xe-x/1/0:[0-3]</code>	<code>et-x/1/0</code>	-

Table 55: Interface Naming Convention for MPC7E-MRATE (Continued)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
	xe-x/1/1:[0-3]	et-x/1/1	-
	xe-x/1/2:[0-3]	et-x/1/2	et-x/1/2
	xe-x/1/3:[0-3]	et-x/1/3	-
	xe-x/1/4:[0-3]	et-x/1/4	-
	xe-x/1/5:[0-3]	et-x/1/5	et-x/1/5

Table 56: Active Physical Ports on MPC7E-MRATE

Number of Active ports Configured	Active Physical Ports for Different Configured Speeds		
	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
1	0	0	2
2	0, 1	0, 1	2, 5
3	0, 1, 2	0, 1, 2	2, 5
4	0, 1, 2, 3	0, 1, 2, 3	2, 5
5	0, 1, 2, 3, 4	0, 1, 2, 3, 4	2, 5
6	0, 1, 2, 3, 4, 5	0, 1, 2, 3, 4, 5	2, 5

For information on how to configure the speed at the PIC level, see ["Configure Port Speed - PIC Level" on page 177](#). For information on how to configure the speed at the port level, see ["Configure Port Speed - Port Level" on page 178](#).

Port Speed on MPC7E-10G

Starting with Junos OS Release 20.4R1, you can configure 1-Gbps speed on 10-Gigabit Ethernet ports of the MPC7E-10G. You can configure the 10-Gigabit Ethernet ports to operate as 1-Gigabit Ethernet port using the speed statement at the `edit interfaces interfacename gigether-options` hierarchy level.

By default, all ports in the MPC7E-10G MPC operate as 10 Gigabit Ethernet interfaces. When you configure the port speed, the PIC and MIC operate at the configured speed. An alarm is generated if you configure an invalid speed and when you change an existing port speed.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MPC7E-10G, see [Hardware Compatibility Tool \(HCT\)](#).

Table 57: Port Speed for MPC7E-10G

PIC	Port	Port Speed Supported
PIC 0	0-19	10-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC 1	0-19	10-Gigabit Ethernet Default: 10-Gigabit Ethernet

[Table 58 on page 227](#) lists the naming conventions used for interfaces on MPC7E-10G for MX240, MX480, MX960, MX2010, and MX2020 routers.

Table 58: Interface Naming Convention for MPC7E-10G

Packet Forwarding Engine	10-Gigabit Ethernet Interface
0	<code>xe-x/0/0</code> to <code>xe-x/0/19</code>
1	<code>xe-x/1/0</code> to <code>xe-x/1/19</code>

Table 59: Active Physical Ports on MPC7E-10G

Number of Active ports Configured	10-Gigabit Ethernet
1	0
2	0, 1
3	0, 1, 2
4	0, 1, 2, 3
5	0, 1, 2, 3, 4
6	0, 1, 2, 3, 4, 5
7	0, 1, 2, 3, 4, 5, 6
8	0, 1, 2, 3, 4, 5, 6, 7
9	0, 1, 2, 3, 4, 5, 6, 7, 8
10	0, 1, 2, 3, 4, 5, 6, 7, 8, 9
11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10
12	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11
13	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12
14	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13
15	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14

Table 59: Active Physical Ports on MPC7E-10G (Continued)

Number of Active ports Configured	10-Gigabit Ethernet
16	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15
17	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16
18	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17
19	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18
20	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port Speed on MIC-MRATE

On a MIC-MRATE MPC by default, all the ports operate as four 10 Gigabit Ethernet interfaces. When you change an existing port speed configuration at the port level for MPC8E or MPC9E, you must reset the MIC using the request chassis mic mic-slot mic-slot-number fpc-slot fpc-slot-number (offline | online) command.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MIC-MRATE, see [Hardware Compatibility Tool \(HCT\)](#).

When you configure the port speed at the PIC level and the port level simultaneously, an error message is displayed, and when you change an existing port speed configuration with an invalid port speed configuration, an alarm is generated.

Table 60: Port Speed for MIC-MRATE

MIC	Port	Port Speed Supported
MIC	12	4x10-Gigabit Ethernet
		10-Gigabit Ethernet
		40-Gigabit Ethernet
		100-Gigabit Ethernet
		Default: 10-Gigabit Ethernet

[Table 61 on page 230](#) lists the naming conventions used for interfaces on MIC-MRATE installed on Slot 0 of MPC8E and MPC9E.

Table 61: Interface Naming Convention for MIC-MRATE

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe-x/0/0:[0-3]-xe-x/0/5:[0-3]	et-x/0/0:[0-3]-xe-x/0/5:[0-3]	et-x/0/0:[0-3]-xe-x/0/3:[0-3]
1	xe-x/0/6:[0-3]-xe-x/0/11:[0-3]	et-x/0/6:[0-3]-xe-x/0/11:[0-3]	et-x/0/6:[0-3]-xe-x/0/11:[0-3]

[Table 62 on page 230](#) lists the naming conventions used for interfaces on MIC-MRATE installed on Slot 1 of MX10003 MPC.

Table 62: Interface Naming Convention for MIC-MRATE

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe-x/0/0:[0-3] to xe-x/0/3:[0-3]	et-x/0/0:[0-3] to xe-x/0/3:[0-3]	et-x/0/0:[0-3] to xe-x/0/3:[0-3]

Table 62: Interface Naming Convention for MIC-MRATE (Continued)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
1	xe-x/0/4:[0-3] to xe-x/0/7:[0-3]	et-x/0/4:[0-3] to xe-x/0/7:[0-3]	et-x/0/6:[0-3] to xe-x/0/7:[0-3]
2	xe-x/0/8:[0-3] to xe-x/0/11:[0-3]	xe-x/0/8:[0-3] to xe-x/0/11:[0-3]	xe-x/0/8:[0-3] to xe-x/0/9:[0-3]

Table 63: Active Physical Ports on MIC-MRATE (MPC8E)

Number of Active Ports Configured	Active Physical Ports for Different Configured Speeds		
	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
1	0	0	0
2	0, 6	0, 6	0, 6
3	0, 1, 6	0, 1, 6	0, 1, 6
4	0, 1, 6, 7	0, 1, 6, 7	0, 1, 6, 7
5	0, 1, 2, 6, 7	0, 1, 2, 6, 7	0, 1, 6, 7
6	0, 1, 2, 6, 7, 8	0, 1, 2, 6, 7, 8	0, 1, 6, 7
7	0, 1, 2, 3, 6, 7, 8	0, 1, 2, 3, 6, 7, 8	0, 1, 6, 7
8	0, 1, 2, 3, 6, 7, 8, 9	0, 1, 2, 3, 6, 7, 8, 9	0, 1, 6, 7
9	0, 1, 2, 3, 4, 6, 7, 8, 9	0, 1, 2, 3, 4, 6, 7, 8, 9	0, 1, 6, 7

Table 63: Active Physical Ports on MIC-MRATE (MPC8E) (Continued)

Number of Active Ports Configured	Active Physical Ports for Different Configured Speeds		
	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
10	0, 1, 2, 3, 4, 6, 7, 8, 9, 10	0, 1, 2, 3, 4, 6, 7, 8, 9, 10	0, 1, 6, 7
11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	0, 1, 6, 7
12	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11	0, 1, 6, 7

[Table 64 on page 232](#) lists the active physical ports on MIC-MRATE (MPC9E and MPC8E (1.6T))

Table 64: Active Physical Ports on MIC-MRATE (MPC9E and MPC8E (1.6T mode))

Number of Active Ports Configured	Active Physical Ports for Different Configured Speeds		
	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
1	0	0	0
2	0, 6	0, 6	0, 6
3	0, 1, 6	0, 1, 6	0, 1, 6
4	0, 1, 6, 7	0, 1, 6, 7	0, 1, 6, 7
5	0, 1, 2, 6, 7	0, 1, 2, 6, 7	0, 1, 2, 6, 7
6	0, 1, 2, 6, 7, 8	0, 1, 2, 6, 7, 8	0, 1, 2, 6, 7, 8
7	0, 1, 2, 3, 6, 7, 8	0, 1, 2, 3, 6, 7, 8	0, 1, 2, 3, 6, 7, 8
8	0, 1, 2, 3, 6, 7, 8, 9	0, 1, 2, 3, 6, 7, 8, 9	0, 1, 2, 3, 6, 7, 8, 9

Table 64: Active Physical Ports on MIC-MRATE (MPC9E and MPC8E (1.6T mode)) (Continued)

Number of Active Ports Configured	Active Physical Ports for Different Configured Speeds		
	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
9	0, 1, 2, 3, 4, 6, 7, 8, 9	0, 1, 2, 3, 4, 6, 7, 8, 9	0, 1, 2, 3, 6, 7, 8, 9
10	0, 1, 2, 3, 4, 6, 7, 8, 9, 10	0, 1, 2, 3, 4, 6, 7, 8, 9, 10	0, 1, 2, 3, 6, 7, 8, 9
11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	0, 1, 2, 3, 6, 7, 8, 9
12	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11	0, 1, 2, 3, 6, 7, 8, 9

For information on how to configure the speed at the PIC level, see ["Configure Port Speed - PIC Level" on page 177](#). For information on how to configure the speed at the port level, see ["Configure Port Speed - Port Level" on page 178](#).

Port Speed on MIC-3D-10GE-SFP-E

The MIC-3D-10GE-SFP-E supports 1G and 10G port speeds through the pic mode where all the ports have the same speed. You can create 10 non-channelized 1GE ports in 1G PIC mode, and 10 non-channelized 10GE ports in 10G PIC.

When you change the mode between 1G and 10G, a pic bounce is triggered. In both modes, the port number will be in the range of 0-9 for the IFDs. When an incorrect PIC mode other than 1G or 10G is committed, the PIC operates on the previously configured PIC mode and an alarm is raised on the incorrect PIC. When an incorrect PIC mode is present while the PIC is booting up, the PIC is configured to the default 1G PIC mode and an alarm is raised.

The default PIC level mode is 1G, the default number of ports is 10G, and by default, auto-negotiation is enabled for 1G interface.

Table 65: Port Speed for MIC-3D-10GE-SFP-E

PIC	Port Number	Port Speed Supported
PIC 0	0-9	1-Gigabit Ethernet 10- Gigabit Ethernet Default: 1-Gigabit Ethernet
PIC 2	0-9	1-Gigabit Ethernet 10- Gigabit Ethernet Default: 1-Gigabit Ethernet

Table 66: Interface Naming Conventions (MIC-3D-10GE-SFP-E)

PIC Mode	Interface Name
1-Gigabit Ethernet Interface	ge-x/[0 2]/[0-9]
10-Gigabit Ethernet Interface	xe-x/[0 2]/[0-9]

For information on how to configure the speed at the PIC level, see ["Configure Port Speed - PIC Level" on page 177](#). For information on how to configure the speed at the port level, see ["Configure Port Speed - Port Level" on page 178](#).

Port Speed on MX10003 MPC

MX10003 router comes up in PIC (PIC 0 and PIC 1) mode and all the interfaces operate at 10-Gbps. To configure a different port speed, switch from PIC mode to port mode and reset both the PICs. You must configure *both the PICs and all the associated ports*, under the [edit chassis] hierarchy.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MX10003 MPC, see [Hardware Compatibility Tool \(HCT\)](#).

When you configure the port speed at the port level, only the configured ports are active. Other ports are disabled. For example, if you want to configure the port speed as 10 Gbps, the port speed of the fixed-port PIC and the MIC should be configured to 10 Gbps. If you want to configure the port speed as 40 Gbps, the port speed of the fixed-port PIC and the MIC should be configured to 40 Gbps. However, if you choose to configure all ports of the MX10003 MPC to operate as 100-Gigabit Ethernet interfaces, the ports on the MIC have to be configured to 100 Gbps and the number-of-ports *number-of-active-physical-ports* statement on the fixed-port PIC must be set to 0.

You cannot configure the ports that will oversubscribe the Packet Forwarding Engine. For example, a combination of eleven 100-Gigabit Ethernet interfaces on the MIC and ten 10-Gigabit Ethernet interfaces on the fixed-port PIC will result in an invalid configuration. If you try to commit an invalid configuration, the configuration will get committed. However, the port will not be activated. You can execute the show chassis alarms to display the error message. The valid configuration in this case would be eleven 100-Gigabit Ethernet interfaces on the MIC and eight 10-Gigabit Ethernet interfaces on the fixed-port PIC.

Table 67: Port Speed for MX10003 MPC

PIC	Port Number	Port Speed Supported
PIC 0 (Fixed-port PIC)	0-5	40-Gigabit Ethernet 4x10-Gigabit Ethernet All the 4x10-Gigabit Ethernet ports can operate in 1 Gbps mode.
PIC 1 (Multi-Rate MIC)	0-11	40-Gigabit Ethernet 4x10-Gigabit Ethernet 100-Gigabit Ethernet On the non-MASSEC MIC, all the 4x10-Gigabit Ethernet ports can operate in 1 Gbps mode.

Table 68: PFE Based Port Mode Configuration

Port Speed configuration on PIC1(Gbps)				Port speed configuration on PIC0(Gbps)	
100	100	100	100	0	0
100	100	100	10/40	0	0

Table 68: PFE Based Port Mode Configuration (Continued)

Port Speed configuration on PIC1(Gbps)				Port speed configuration on PIC0(Gbps)	
100	100	100	0	10/40	10/40
100	100	10/40	10/40	10/40	10/40
100	10/40	10/40	10/40	10/40	10/40
10/40	10/40	10/40	0	10/40	10/40
10/40	10/40	10/40	10/40	0	0

 **NOTE:** Use the [port-checker tool](#) to check whether the combination of ports you want to use is valid or not.

Table 69: PIC Mode Configuration

Port Speed configuration on PIC1(Gbps)		Port speed configuration on PIC0(Gbps)
100		0 Configure the number of ports to 0.
10		10
40		40

[Table 70 on page 237](#) lists the naming conventions used for interfaces on the fixed-port PIC when installed in slot 0 of the MX10003 MPC. [Table 71 on page 237](#) lists the naming conventions used for interfaces on the modular MIC when installed in slot 1 of the MPC.

Table 70: Interface Naming Convention for the Fixed-Port PIC in Slot 0 of MX10003 MPC

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe-x/0/0:[0-3]	et-x/0/0	-
	xe-x/0/1:[0-3]	et-x/0/1	-
1	xe-x/0/2:[0-3]	et-x/0/2	-
	xe-x/0/3:[0-3]	et-x/0/3	-
2	xe-x/0/4:[0-3]	et-x/0/4	-
	xe-x/0/5:[0-3]	et-x/0/5	-

Table 71: Interface Naming Convention for Modular MIC Installed in Slot 1 of MX10003 MPC

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe-x/1/0:[0-3]	et-x/1/0	et-x/1/0
	xe-x/1/1:[0-3]	et-x/1/1	et-x/1/1
	xe-x/1/2:[0-3]	et-x/1/2	et-x/1/2
	xe-x/1/3:[0-3]	et-x/1/3	et-x/1/3
1	xe-x/1/4:[0-3]	et-x/1/4	et-x/1/4
	xe-x/1/5:[0-3]	et-x/1/5	et-x/1/5

Table 71: Interface Naming Convention for Modular MIC Installed in Slot 1 of MX10003 MPC
(Continued)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
	xe-x/1/6:[0-3]	et-x/1/6	et-x/1/6
	xe-x/1/7:[0-3]	et-x/1/7	et-x/1/7
2	xe-x/1/8:[0-3]	et-x/1/8	et-x/1/8
	xe-x/1/9:[0-3]	et-x/1/9	et-x/1/9
	xe-x/1/10:[0-3]	et-x/1/10	et-x/1/10
	xe-x/1/11:[0-3]	et-x/1/11	et-x/1/11

Table 72: Active Physical Ports on the MX10003 MPC at the MIC level

Ports Configured (number-of-ports Statement)	Active Physical Ports for Different Configured Speeds		
	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
1	0	0	0
2	0, 4	0, 4	0, 4
3	0, 4, 8	0, 4, 8	0, 4, 8
4	0, 1, 4, 8	0, 1, 4, 8	0, 1, 4, 8
5	0, 1, 4, 5, 8	0, 1, 4, 5, 8	0, 1, 4, 5, 8

Table 72: Active Physical Ports on the MX10003 MPC at the MIC level (Continued)

Ports Configured (number-of-ports Statement)	Active Physical Ports for Different Configured Speeds		
	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
6	0, 1, 4, 5, 8, 9	0, 1, 4, 5, 8, 9	0, 1, 4, 5, 8, 9
7	0, 1, 2, 4, 5, 8, 9	0, 1, 2, 4, 5, 8, 9	0, 1, 2, 4, 5, 8, 9
8	0, 1, 2, 4, 5, 6, 8, 9	0, 1, 2, 4, 5, 6, 8, 9	0, 1, 2, 4, 5, 6, 8, 9
9	0, 1, 2, 4, 5, 6, 8, 9, 10	0, 1, 2, 4, 5, 6, 8, 9, 10	0, 1, 2, 4, 5, 6, 8, 9, 10
10	0, 1, 2, 3, 4, 5, 6, 8, 9, 10	0, 1, 2, 3, 4, 5, 6, 8, 9, 10	0, 1, 2, 3, 4, 5, 6, 8, 9, 10
11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10
12	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11

Table 73: Active Physical Ports on MX10003 MPC at the PIC level

Ports Configured (number-of-ports Statement)	Active Physical Ports for Different Configured Speeds	
	10-Gigabit Ethernet	40-Gigabit Ethernet
1	0	0
2	0, 2	0, 2
3	0, 2, 4	0, 2, 4

Table 73: Active Physical Ports on MX10003 MPC at the PIC level (Continued)

Ports Configured (number-of-ports Statement)	Active Physical Ports for Different Configured Speeds	
	10-Gigabit Ethernet	40-Gigabit Ethernet
4	0, 1, 2, 4	0, 1, 2, 4
5	0, 1, 2, 3, 4	0, 1, 2, 3, 4
6	0, 1, 2, 3, 4, 5	0, 1, 2, 3, 4, 5

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port Speed on MX10K-LC4800

A PIC port in MX10K-LC4800 support multiple port speeds. A significant number of port speeds requires breakout optics. MX10K-LC4800 has multiple variants of a breakout optics. Hence, you require a default port channelization to have a deterministic system behavior.

For information about the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs, see [Hardware Compatibility Tool \(HCT\)](#).

Table 74: Port Speed for MX10K-LC4800

PIC	Port Number	Port Speed Supported
PIC 0 (12xSFP56-DD + 2xQSFP56-DD)	0-1	4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 2x100-Gigabit Ethernet 4x100-Gigabit Ethernet 400-Gigabit Ethernet
	2-13	1-Gigabit Ethernet 10-Gigabit Ethernet 25-Gigabit Ethernet 50-Gigabit Ethernet 100-Gigabit Ethernet
PIC 1 (12xSFP56-DD + 2xQSFP56-DD)	0-1	4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 2x100-Gigabit Ethernet 4x100-Gigabit Ethernet 400-Gigabit Ethernet

Table 74: Port Speed for MX10K-LC4800 (*Continued*)

PIC	Port Number	Port Speed Supported
	2-13	1-Gigabit Ethernet 10-Gigabit Ethernet 25-Gigabit Ethernet 50-Gigabit Ethernet 100-Gigabit Ethernet
PIC 2 (16xSFP56-DD)	0-15	1-Gigabit Ethernet 10-Gigabit Ethernet 25-Gigabit Ethernet 50-Gigabit Ethernet 100-Gigabit Ethernet

Table 75: Interface Naming Convention for MX10K-LC4800

Interface Type	Interfaces	Notes
1-Gigabit Ethernet Interface	ge-x/y/z	
10-Gigabit Ethernet Interface	xe-x/y/z	
4x10-Gigabit Ethernet Interface	xe-x/y/z:0 xe-x/y/z:1 xe-x/y/z:2 xe-x/y/z:3	<ul style="list-style-type: none"> • x refers to the FPC slot number. • y refers to the PIC slot number and the valid range is 0 - 2. • z refers to the physical port number and the valid range is 0 - 13 or 0 - 15 for LC4800.
25-Gigabit Ethernet Interface	et-x/y/z	

Table 75: Interface Naming Convention for MX10K-LC4800 (Continued)

Interface Type	Interfaces	Notes
4x25-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1 et-x/y/z:2 et-x/y/z:3	
40-Gigabit Ethernet Interface	et-x/y/z	
50-Gigabit Ethernet Interface	et-x/y/z	
100-Gigabit Ethernet Interface	et-x/y/z	
2x100-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1	
4x100-Gigabit Ethernet Interface	et-x/y/z:0 et-x/y/z:1 et-x/y/z:2 et-x/y/z:3	
400-Gigabit Ethernet Interface	et-x/y/z	

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port Speed on MIC-MACSEC-20GE

On a MIC-MACSEC-20GE MIC, all the ports operate as 1-Gigabit Ethernet interface, and the ports in the MIC are created as `ge` interfaces distributed across PIC0 and PIC1. In 10-Gigabit Ethernet mode, the ports in the MIC will be created as `xe` interfaces one each on PIC 0 and PIC 1.

The MIC-MACSEC-20GE MIC provides 128-bit and 256-bit MACsec encryption on all the twenty 1-Gigabit Ethernet and the two 10-Gigabit Ethernet ports. By default, 128-bit MACsec encryption is supported.

Configure both the PICs in the `pic-mode`, to set the PIC to operate in 10G mode, use the `pic-mode 10G` command. For any misconfiguration, the PIC operates in default `pic-mode` (20x1GE), where all ports are in 1-GE port speed.

For information on the MIC, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MIC-MACSEC-20GE, see [Hardware Compatibility Tool \(HCT\)](#).

When you change the speed of the port from 1-Gbps to 10-Gbps or vice versa, the Flexible PIC Concentrator (FPC) on MX240, MX480, and MX960 routers, and the Forwarding Engine Board (FEB) in MX80 and MX104 routers reboot. The 10-Gbps ports (ports 8 and 9) of the MIC-MACSEC-20GE MIC display the link status as up while the peer side is down. In this case, it is recommended to disable auto-negotiation and set the speed to 1-Gbps on the peer side to bring the link up on the peer side.

Table 76: Port Speed for MIC-MACSEC-20GE

PIC	Port Number	Port Speed Supported
PIC 0	0-9	1-Gigabit Ethernet Default: 1-Gigabit Ethernet
PIC 1	0-7	1-Gigabit Ethernet Default: 1-Gigabit Ethernet
PIC 1	8-9	10-Gigabit Ethernet Default: 10-Gigabit Ethernet

Table 77: Interface Naming Conventions (MIC-MACSEC-20GE)

PIC	1-Gigabit Ethernet Interface	10-Gigabit Ethernet Interface
PIC 0	ge-x/0/[0-9]	xe-x/0/0
PIC 1	ge-x/1/[0-9]	xe-x/1/0
PIC 2	ge-x/2/[0-9]	xe-x/2/0
PIC 3	ge-x/3/[0-3]	xe-x/3/0

For information on how to configure the speed at the PIC level, see ["Configure Port Speed - PIC Level" on page 177](#). For information on how to configure the speed at the port level, see ["Configure Port Speed - Port Level" on page 178](#).

Port Speed on MPC10E-10C-MRATE

On an MPC10E-10C-MRATE, all the interfaces in the PIC mode operate at 100-Gbps speed on all the PICs (PIC 0, PIC1, and PIC 2). When you change the speed at the PIC level, the PIC restarts automatically with the newly configured speed. When you configure the speed of a port at the port level, only the speed of that port is modified. All other ports in the PIC remain unaffected. You cannot configure the port speed simultaneously at the PIC and port levels.

Oversubscription of Packet Forwarding Engine (PFE) capacity is not supported. The demand on each PFE must be less than or equal to its forwarding capacity. You can configure active ports to prevent oversubscription.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MPC10E-10C-MRATE, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 78 on page 246](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

Table 78: Port Speed for MPC10E-10C-MRATE

PIC	Port Number	Port Speed Supported
PIC 0 (or PFE 0)	0-4	40-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 100-Gigabit Ethernet 400-Gigabit Ethernet (Only on port 4) Default: All active ports operate in 100 Gbps mode.
PIC 1 (or PFE 1)	0-4	40-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 100-Gigabit Ethernet 400-Gigabit Ethernet (Only on port 4) Default: All active ports operate in 100 Gbps mode.

Different PICs in the MPC10E-10C-MRATE can operate at different speeds. That is, PIC speed of one PIC does not apply to the other PICs in the MPC.

The default port speed is 100 Gbps, and you can configure ports without specifying speeds under the set chassis hierarchy. However, if you configure a port's speed under this hierarchy, you must configure all utilized ports on that PIC with the relevant supported speed. Failure to do so will result in only the configured port remaining active, while the others will be deleted. This ensures consistent speed configuration across all ports on the PIC.

Table 79: Port speed support on MPC10E-10C-MRATE

	PIC level Profile					Port level Profile
	10G	25G	40G	100G	400G	
PIC 0 (5xQSFP28 PIC)	Yes	Yes	Yes	Yes	Yes	Yes
PIC 1 (5xQSFP28 PIC)	Yes	Yes	Yes	Yes	Yes	Yes

Table 80: Port Speed Combinations with 400GE on MPC10E-10C-MRATE

PIC	Port Speed Combinations	Supported
PIC 0	400GE + 4x10GE	Yes
	400GE + 4x25GE	Yes
	400GE + 40GE	Yes
	400GE + 100GE	Yes
	400GE + 8x10GE (using 2x 4x10GE breakout cables)	No
	400GE + 8x25GE (using 2x 4x25GE breakout cables)	No
	400GE + 2x40GE	No
PIC 1	400GE + 4x10GE	Yes
	400GE + 4x25GE	Yes
	400GE + 40GE	Yes

Table 80: Port Speed Combinations with 400GE on MPC10E-10C-MRATE (Continued)

PIC	Port Speed Combinations	Supported
	400GE + 100GE	Yes
	400GE + 8x10GE (using 2x 4x10GE breakout cables)	No
	400GE + 8x25GE (using 2x 4x25GE breakout cables)	No
	400GE + 2x40GE	No

Table 81: Interface Naming Convention for MPC10E-10C-MRATE

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
0	xe-x/0/0:[0-3]	et-x/0/0:[0-3]	et-x/0/0	et-x/0/0	-
	xe-x/0/1:[0-3]	et-x/0/1:[0-3]	et-x/0/1	et-x/0/1	-
	xe-x/0/2:[0-3]	et-x/0/2:[0-3]	et-x/0/2	et-x/0/2	-
	xe-x/0/3:[0-3]	et-x/0/3:[0-3]	et-x/0/3	et-x/0/3	-
	xe-x/0/4:[0-3]	et-x/0/4:[0-3]	et-x/0/4	et-x/0/4	et-x/0/4
1	xe-x/1/0:[0-3]	et-x/1/0:[0-3]	et-x/1/0	et-x/1/0	-
	xe-x/1/1:[0-3]	et-x/1/1:[0-3]	et-x/1/1	et-x/1/1	-
	xe-x/1/2:[0-3]	et-x/1/2:[0-3]	et-x/1/2	et-x/1/2	-

Table 81: Interface Naming Convention for MPC10E-10C-MRATE (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	xe-x/1/3:[0-3]	et-x/1/3:[0-3]	et-x/1/3	et-x/1/3	-
	xe-x/1/4:[0-3]	et-x/1/4:[0-3]	et-x/1/4	et-x/1/4	et-x/0/4

Table 82: Active Ports with port speed configured at PIC level

PIC Type	Number of Active Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSF P28 PIC (PIC 0)	1	0	0	0	0	4
	2	0, 1	0,1	0, 1	0, 1	4
	3	0, 1, 2	0,1,2	0, 1, 2	0, 1, 2	4
	4	0, 1, 2, 3	0,1,2,3	0, 1, 2, 3	0, 1, 2, 3	4
	5	0, 1, 2, 3, 4	0,1,2,3,4	0, 1, 2, 3, 4	0, 1, 2, 3, 4	4
5xQSF P28 PIC (PIC 1)	1	0	0	0	0	4
	2	0, 1	0,1	0, 1	0, 1	4
	3	0, 1, 2	0,1,2	0, 1, 2	0, 1, 2	4
	4	0, 1, 2, 3	0,1,2,3	0, 1, 2, 3	0, 1, 2, 3	4

Table 82: Active Ports with port speed configured at PIC level (*Continued*)

PIC Type	Number of Active Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
	5	0, 1, 2, 3, 4	0,1,2,3,4	0, 1, 2, 3, 4	0, 1, 2, 3, 4	4

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port Speed on MPC10E-15C-MRATE

On an MPC10E-15C-MRATE, all the interfaces in the PIC mode operate at 100-Gbps speed on all the PICs (PIC 0, PIC1, and PIC 2). When you change the speed at the PIC level, the PIC restarts automatically with the newly configured speed. When you configure the speed of a port at the port level, only the speed of that port is modified. All other ports in the PIC remain unaffected. You cannot configure the port speed simultaneously at the PIC and port levels.

Oversubscription of Packet Forwarding Engine (PFE) capacity is not supported. The demand on each PFE must be less than or equal to its forwarding capacity. You can configure active ports to prevent oversubscription.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MPC10E-15C-MRATE, see [Hardware Compatibility Tool \(HCT\)](#).

Table 83: Port Speed for MPC10E-15C-MRATE

PIC	Port Number	Port Speed Supported
PIC 0 (or PFE 0)	0-4	40-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 100-Gigabit Ethernet 400-Gigabit Ethernet (Only on port 4) Default: All active ports operate in 100 Gbps mode.
PIC 1 (or PFE 1)	0-4	40-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 100-Gigabit Ethernet 400-Gigabit Ethernet (Only on Port 4) Default: All active ports operate in 100 Gbps mode.
PIC 2 (or PFE 2)	0-4	40-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 100-Gigabit Ethernet 400-Gigabit Ethernet (Only on port 4) Default: All active ports operate in 100 Gbps mode.

Different PICs in the MPC10E-15C-MRATE can operate at different speeds. That is, PIC speed of one PIC does not apply to the other PICs in the MPC.

The default port speed is 100 Gbps, and you can configure ports without specifying speeds under the set chassis hierarchy. However, if you configure a port's speed under this hierarchy, you must configure all utilized ports on that PIC with the relevant supported speed. Failure to do so will result in only the

configured port remaining active, while the others will be deleted. This ensures consistent speed configuration across all ports on the PIC.

Table 84: Port speed support on MPC10E-15C-MRATE

PIC	PIC level Profile					Port level Profile
	10G	25G	40G	100G	400G	
PIC 0 (5xQSFP28 PIC)	Yes	Yes	Yes	Yes	Yes	Yes
PIC 1 (5xQSFP28 PIC)	Yes	Yes	Yes	Yes	Yes	Yes
PIC 2 (5xQSFP28 PIC)	Yes	Yes	Yes	Yes	Yes	Yes

Table 85: Port Speed Combinations with 400GE on MPC10E-15C-MRATE

PIC	Port Speed Combinations	Supported
PIC 0	400GE + 4x10GE	Yes
	400GE + 4x25GE	Yes
	400GE + 40GE	Yes
	400GE + 100GE	Yes
	400GE + 8x10GE (using 2x 4x10GE breakout cables)	No
	400GE + 8x25GE (using 2x 4x25GE breakout cables)	No
	400GE + 2x40GE	No
PIC 1	400GE + 4x10GE	Yes

Table 85: Port Speed Combinations with 400GE on MPC10E-15C-MRATE (Continued)

PIC	Port Speed Combinations	Supported
	400GE + 4x25GE	Yes
	400GE + 40GE	Yes
	400GE + 100GE	Yes
	400GE + 8x10GE (using 2x 4x10GE breakout cables)	No
	400GE + 8x25GE (using 2x 4x25GE breakout cables)	No
	400GE + 2x40GE	No
PIC 2	400GE + 4x10GE	Yes
	400GE + 4x25GE	Yes
	400GE + 40GE	Yes
	400GE + 100GE	Yes
	400GE + 8x10GE (using 2x 4x10GE breakout cables)	No
	400GE + 8x25GE (using 2x 4x25GE breakout cables)	No
	400GE + 2x40GE	No

Use the number-of-sub-ports *<number-of-sub-ports>* configuration statement to specify the number of interfaces to be created for a specified speed. You can configure the *number-of-sub-ports* only for 10-

Gbps speed, set the *number-of-sub-ports* to 4, to channelize 40-Gigabit Ethernet interface to four 10-Gigabit Ethernet interfaces.

Table 86: Interface Naming Convention for MPC10E-15C-MRATE

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
0	xe-x/0/0:[0-3]	et-x/0/0:[0-3]	et-x/0/0	et-x/0/0	-
	xe-x/0/1:[0-3]	et-x/0/1:[0-3]	et-x/0/1	et-x/0/1	-
	xe-x/0/2:[0-3]	et-x/0/2:[0-3]	et-x/0/2	et-x/0/2	-
	xe-x/0/3:[0-3]	et-x/0/3:[0-3]	et-x/0/3	et-x/0/3	-
	xe-x/0/4:[0-3]	et-x/0/4:[0-3]	et-x/0/4	et-x/0/4	et-x/0/4
1	xe-x/1/0:[0-3]	et-x/1/0:[0-3]	et-x/1/0	et-x/1/0	-
	xe-x/1/1:[0-3]	et-x/1/1:[0-3]	et-x/1/1	et-x/1/1	-
	xe-x/1/2:[0-3]	et-x/1/2:[0-3]	et-x/1/2	et-x/1/2	-
	xe-x/1/3:[0-3]	et-x/1/3:[0-3]	et-x/1/3	et-x/1/3	-
	xe-x/1/4:[0-3]	et-x/1/4:[0-3]	et-x/1/4	et-x/1/4	et-x/1/4
2	xe-x/2/0:[0-3]	et-x/2/0:[0-3]	et-x/2/0	et-x/2/0	-
	xe-x/2/1:[0-3]	et-x/2/1:[0-3]	et-x/2/1	et-x/2/1	-
	xe-x/2/2:[0-3]	et-x/2/2:[0-3]	et-x/2/2	et-x/2/2	-

Table 86: Interface Naming Convention for MPC10E-15C-MRATE (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	xe-x/2/3:[0-3]	et-x/2/3: [0-3]	et-x/2/3	et-x/2/3	-
	xe-x/2/4:[0-3]	et-x/2/4: [0-3]	et-x/2/4	et-x/2/4	et-x/2/4

Table 87: Active Ports with port speed configuration at PIC level

PIC Type	Number of Active Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSFP 28 PIC (PIC 0)	1	0	0	0	0	4
	2	0, 1	0,1	0, 1	0, 1	4
	3	0, 1, 2	0,1,2	0, 1, 2	0, 1, 2	4
	4	0, 1, 2, 3	0,1,2,3	0, 1, 2, 3	0, 1, 2, 3	4
	5	0, 1, 2, 3, 4	0,1,2,3,4	0, 1, 2, 3, 4	0, 1, 2, 3, 4	4
5xQSFP 28 PIC (PIC 1)	1	0	0	0	0	4
	2	0, 1	0,1	0, 1	0, 1	4
	3	0, 1, 2	0,1,2	0, 1, 2	0, 1, 2	4
	4	0, 1, 2, 3	0,1,2,3	0, 1, 2, 3	0, 1, 2, 3	4
	5	0, 1, 2, 3, 4	0,1,2,3,4	0, 1, 2, 3, 4	0, 1, 2, 3, 4	4

Table 87: Active Ports with port speed configuration at PIC level (Continued)

PIC Type	Number of Active Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSFP 28 PIC (PIC 2)	1	0	0	0	0	4
	2	0, 1	0,1	0, 1	0, 1	4
	3	0, 1, 2	0,1,2	0, 1, 2	0, 1, 2	4
	4	0, 1, 2, 3	0,1,2,3	0, 1, 2, 3	0, 1, 2, 3	4
	5	0, 1, 2, 3, 4	0,1,2,3,4	0, 1, 2, 3, 4	0, 1, 2, 3, 4	4

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port Speed on MX2K-MPC11E

MX2K-MPC11E MPC supports a speed of 100 Gbps by default on ports 0 through port 4 on each PIC. When you change the speed at the PIC level, the PIC restarts automatically with the newly configured speed. Similarly, when you change the port configuration at the port level, the interfaces corresponding to the affected port are deleted and later recreated.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the supported transceivers, optical interfaces, and DAC cables on MX2K-MPC11E, see [Hardware Compatibility Tool \(HCT\)](#).

When you configure the port speed at the PIC level and the port level simultaneously, an error message is displayed, and when you change an existing port speed configuration with an invalid port speed configuration, an alarm is generated.

You can configure a maximum bandwidth of 500 Gbps per PIC. Restart the PIC manually each time you configure the port speed of a port on any PIC.

Table 88: Port Speed for the MX2K-MPC11E

PIC	Port Number	Port Speed Supported
PIC 0 (or PFE 0)	0	40-Gigabit Ethernet
		4x10-Gigabit Ethernet
	1-4	100-Gigabit Ethernet
		400-Gigabit Ethernet
	0	4x100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
PIC 1 (or PFE 1)	1-4	100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	0	40-Gigabit Ethernet
		4x10-Gigabit Ethernet
	1-4	100-Gigabit Ethernet
		400-Gigabit Ethernet
	0	4x100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	1-4	100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.

Table 88: Port Speed for the MX2K-MPC11E *(Continued)*

PIC	Port Number	Port Speed Supported
PIC 2 (or PFE 2)	0	40-Gigabit Ethernet 4x10-Gigabit Ethernet 100-Gigabit Ethernet 400-Gigabit Ethernet 4x100-Gigabit Ethernet Default: All active ports operate in 100-Gigabit Ethernet mode.
	1-4	100-Gigabit Ethernet Default: All active ports operate in 100-Gigabit Ethernet mode.
PIC 3 (or PFE 3)	0	40-Gigabit Ethernet 4x10-Gigabit Ethernet 100-Gigabit Ethernet 400-Gigabit Ethernet 4x100-Gigabit Ethernet Default: All active ports operate in 100-Gigabit Ethernet mode.
	1-4	100-Gigabit Ethernet Default: All active ports operate in 100-Gigabit Ethernet mode.

Table 88: Port Speed for the MX2K-MPC11E *(Continued)*

PIC	Port Number	Port Speed Supported
PIC 4 (or PFE 4)	0	40-Gigabit Ethernet
		4x10-Gigabit Ethernet
PIC 5 (or PFE 5)	0	100-Gigabit Ethernet
		400-Gigabit Ethernet
	1-4	4x100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	1-4	100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	0	40-Gigabit Ethernet
		4x10-Gigabit Ethernet
	1-4	100-Gigabit Ethernet
		400-Gigabit Ethernet
	0	4x100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	1-4	100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.

Table 88: Port Speed for the MX2K-MPC11E *(Continued)*

PIC	Port Number	Port Speed Supported
PIC 6 (or PFE 6)	0	40-Gigabit Ethernet
		4x10-Gigabit Ethernet
PIC 7 (or PFE 7)	0	100-Gigabit Ethernet
		400-Gigabit Ethernet
	1-4	4x100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	1-4	100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	0	40-Gigabit Ethernet
		4x10-Gigabit Ethernet
	1-4	100-Gigabit Ethernet
		400-Gigabit Ethernet
	0	4x100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.
	1-4	100-Gigabit Ethernet
		Default: All active ports operate in 100-Gigabit Ethernet mode.

Table 89: Interface Naming Convention for MX2K-MPC11E

PIC	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	4x100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
0	xe-x/0/0: [0-3]	et-x/0/0	et-x/0/0 et-x/0/1 et-x/0/2 et-x/0/3 et-x/0/4	et-x/0/0:[0-3]	et-x/0/0 - - - -
1	xe-x/1/0: [0-3]	et-x/1/0	et-x/1/0 et-x/1/1 et-x/1/2 et-x/1/3 et-x/1/4	et-x/1/0:[0-3]	et-x/1/0 - - - -
2	xe-x/2/0: [0-3]	et-x/2/0	et-x/2/0 et-x/2/1 et-x/2/2 et-x/2/3	et-x/2/0:[0-3]	et-x/2/0 - - -

Table 89: Interface Naming Convention for MX2K-MPC11E (Continued)

PIC	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	4x100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
			et-x/2/4		-
3	xe-x/3/0: [0-3]	et-x/3/0	et-x/3/0 et-x/3/1 et-x/3/2 et-x/3/3 et-x/3/4	et-x/3/0:[0-3]	et-x/3/0 - - - -
4	xe-x/4/0: [0-3]	et-x/4/0	et-x/4/0 et-x/4/1 et-x/4/2 et-x/4/3 et-x/4/4	et-x/4/0:[0-3]	et-x/4/0 - - - -
5	xe-x/5/0: [0-3]	et-x/5/0	et-x/5/0 et-x/5/1 et-x/5/2	et-x/5/0:[0-3]	et-x/5/0 - -

Table 89: Interface Naming Convention for MX2K-MPC11E (Continued)

PIC	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	4x100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
			et-x/5/3		-
			et-x/5/4		-
6	xe-x/6/0: [0-3]	et-x/6/0	et-x/6/0	et-x/6/0:[0-3]	et-x/6/0
			et-x/6/1		-
			et-x/6/2		-
			et-x/6/3		-
			et-x/6/4		-
7	xe-x/7/0: [0-3]	et-x/7/0	et-x/7/0	et-x/7/0:[0-3]	et-x/7/0
			et-x/7/1		-
			et-x/7/2		-
			et-x/7/3		-
			et-x/7/4		-

Table 90: Active Ports with port speed configuration at PIC level

PIC Type	Number of Active Ports	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSFP28 PIC (PIC 0)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0
5xQSFP28 PIC (PIC 1)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0
5xQSFP28 PIC (PIC 2)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0

Table 90: Active Ports with port speed configuration at PIC level (Continued)

PIC Type	Number of Active Ports	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSFP28 PIC (PIC 3)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0
5xQSFP28 PIC (PIC 4)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0
5xQSFP28 PIC (PIC 5)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0

Table 90: Active Ports with port speed configuration at PIC level (Continued)

PIC Type	Number of Active Ports	10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSFP28 PIC (PIC 6)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0
5xQSFP28 PIC (PIC 7)	1	0	0	0	0
	2	0	0	0, 1	0
	3	0	0	0, 1, 2	0
	4	0	0	0, 1, 2, 3	0
	5	0	0	0, 1, 2, 3, 4	0

For information on how to configure the speed at the PIC level, see ["Configure Port Speed - PIC Level" on page 177](#). For information on how to configure the speed at the port level, see ["Configure Port Speed - Port Level" on page 178](#).

Port Speed on MX10K-LC480

For information on the line card, see [Protocol and App Support for MX10K-LC480](#)

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#)

For details on software feature support, see [No Link Title](#).

Table 91: MX10K-LC480 line card - Software Support and Description

Software Feature Support	Description
Maximum Transmission Unit (MTU) Size	Maximum MTU of size 16000 bytes for transit traffic. Ingress interface limits is 9000 bytes.
Port Speed Configuration	For Junos OS 21.2R1, you can configure the port speed at the [edit chassis] hierarchy.

Table 92: Port speed for MX10K-LC480

PIC	Port Number	Port Speed Supported
PIC 0	0-23	10-Gigabit Ethernet 1-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC 1	0-23	10-Gigabit Ethernet 1-Gigabit Ethernet Default: 10-Gigabit Ethernet

For information on how to configure the speed at the PIC level, see [Table on page 295](#).

Table 93: Interface Naming Convention for MX10K-LC480

PIC	Interface type	Interfaces
PIC 0	1-Gigabit Ethernet interface (24 SFP+ or SFP ports)	ge-0/0/0 - ge-0/0/23
	10-Gigabit Ethernet interface (24 SFP+ or SFP ports)	xe-0/0/0 - xe-0/0/23

Table 93: Interface Naming Convention for MX10K-LC480 (Continued)

PIC	Interface type	Interfaces
PIC 1	10-Gigabit Ethernet interface (24 SFP+ or SFP ports)	ge-0/1/0 - ge-0/1/23
	10-Gigabit Ethernet interface (24 SFP+ or SFP ports)	xe-0/1/0 - xe-0/1/23

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port Speed on MX10K-LC2101

MX10K-LC2101 MPC operates at the speed of 10 Gbps (using breakout cables), 40 Gbps, and 100 Gbps with a default speed of 10 Gbps. MX10K-LC2101 MPC does not support bandwidth oversubscription hence, when you configure the ports on all PICs, ensure the PFE is less than or equal to its forwarding capacity.

To change the default speed, you must select a port and configure a different port speed on it. For this configuration to take effect, you do not need to reload the Line card.

Starting with Junos OS Release 19.4R1, you can now configure 1-Gbps speed on 10-Gigabit Ethernet ports of the MX10K-LC2101 MPC.

For information on the line card, see [MX10K-LC2101 for MX10008 Routers](#).

To view the supported transceivers, optical interfaces, and DAC cables on MX10K-LC2101, see [Hardware Compatibility Tool \(HCT\)](#).

When you configure the port speed at the PIC level and the port level simultaneously, an error message is displayed, and when you change an existing port speed configuration with an invalid port speed configuration, an alarm is generated.

Each of the 40-Gigabit Ethernet port can be split to four 10-Gigabit Ethernet ports that can be configured to operate as 1-Gigabit Ethernet port. You must use 4x10GE LR breakout optics (QSFP-4X10GE-LR) at the MX10008 or MX10016 end and 1-Gigabit Ethernet EX optics at the remote end. It is only optional to use Juniper optics (SFP-GE40KM) at the remote end, as any vendor's EX (not SX or LX) optics can be used.

To configure the operating speed of the 10-Gbps port to 1-Gbps, use the speed option, the output does not display support for auto-negotiation. However, autonegotiation is supported when the interface speed is configured for 1-Gbps speed.

Table 94: Port Speed for MX10K-LC2101

PIC	Ports	Port Speed Supported
PIC0	0-4	10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC1	0-4	10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC2	0-4	10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC3	0-4	10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC4	0-4	10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet

Table 94: Port Speed for MX10K-LC2101 (Continued)

PIC	Ports	Port Speed Supported
PIC5	0-4	10-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet

When you change the number of active ports using the `number-of-ports` command, you must reset the PIC. Interfaces are created only for active ports. Only the ports you configure are known as the active ports. The number of active ports enables you to handle bandwidth oversubscription. You cannot configure the number of active ports at the port level. If you attempt to do so, an error message is displayed.

Table 95: Interface Naming Conventions (MX10K-LC2101)

PFE	10-Gigabit Ethernet Interface	40-Gigabit Ethernet	100-Gigabit Ethernet
0	<code>xe-x/0/0:[0-3]</code>	<code>et-x/0/0</code>	<code>et-x/0/0</code>
	<code>xe-x/0/1:[0-3]</code>	<code>et-x/0/1</code>	<code>et-x/0/1</code>
	<code>xe-x/0/2:[0-3]</code>	<code>et-x/0/2</code>	<code>et-x/0/2</code>
	<code>xe-x/0/3:[0-3]</code>	<code>et-x/0/3</code>	<code>et-x/0/3</code>
1	<code>xe-x/1/0:[0-3]</code>	<code>et-x/1/0</code>	<code>et-x/1/0</code>
	<code>xe-x/1/1:[0-3]</code>	<code>et-x/1/1</code>	<code>et-x/1/1</code>
	<code>xe-x/1/2:[0-3]</code>	<code>et-x/1/2</code>	<code>et-x/1/2</code>
	<code>xe-x/1/3:[0-3]</code>	<code>et-x/1/3</code>	<code>et-x/1/3</code>

Table 95: Interface Naming Conventions (MX10K-LC2101) (Continued)

PFE	10-Gigabit Ethernet Interface	40-Gigabit Ethernet	100-Gigabit Ethernet
2	xe-x/2/0:[0-3]	et-x/2/0	et-x/2/0
	xe-x/2/1:[0-3]	et-x/2/1	et-x/2/1
	xe-x/2/2:[0-3]	et-x/2/2	et-x/2/2
	xe-x/2/3:[0-3]	et-x/2/3	et-x/2/3
3	xe-x/3/0:[0-3]	et-x/3/0	et-x/3/0
	xe-x/3/1:[0-3]	et-x/3/1	et-x/3/1
	xe-x/3/2:[0-3]	et-x/3/2	et-x/3/2
	xe-x/3/3:[0-3]	et-x/3/3	et-x/3/3
4	xe-x/4/0:[0-3]	et-x/4/0	et-x/4/0
	xe-x/4/1:[0-3]	et-x/4/1	et-x/4/1
	xe-x/4/2:[0-3]	et-x/4/2	et-x/4/2
	xe-x/4/3:[0-3]	et-x/4/3	et-x/4/3
5	xe-x/5/0:[0-3]	et-x/5/0	et-x/5/0
	xe-x/5/1:[0-3]	et-x/5/1	et-x/5/1
	xe-x/5/2:[0-3]	et-x/5/2	et-x/5/2

Table 95: Interface Naming Conventions (MX10K-LC2101) (Continued)

PFE	10-Gigabit Ethernet Interface	40-Gigabit Ethernet	100-Gigabit Ethernet
	xe-x/5/3:[0-3]	et-x/5/3	et-x/5/3

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

Port Speed on MX10K-LC9600

MX10K-LC9600 supports a speed of 400 Gbps by default. You can configure all the four ports as 400 Gbps at the same time.

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs, see [Hardware Compatibility Tool \(HCT\)](#).

For information on MX10K-LC9600, see [Protocols and Application supported by the MX10K-LC9600](#).

When you configure the port speed at the PIC level and the port level simultaneously, an error message is displayed, and when you change an existing port speed configuration with an invalid port speed configuration, an alarm is generated.

Table 96: MX10K-LC9600 - Software Support and Description

Software Feature Support	Description
Maximum Transmission Unit (MTU) Size	Maximum MTU of size 16000 bytes for transit traffic. Minimum interface MTU limits is 274 bytes.
Port Speed Configuration	For Junos OS 21.4R1, you can configure the port speed at the [edit chassis] hierarchy.

Table 96: MX10K-LC9600 - Software Support and Description (*Continued*)

Software Feature Support	Description
Forward Error Correction (FEC) support	<p>By default, 200-Gigabit Ethernet interfaces and 400-Gigabit Ethernet interfaces support FEC119 (KP4) mode.</p> <p>FEC mode is based on the type of optics connected. With non Juniper optics, you can enable FEC mode by manually configuring it. When you manually configure the FEC mode, all the interfaces on a port must be configured with same FEC mode. For example, for a 4x100-Gigabit Ethernet configuration mode, each of the four 100-Gigabit Ethernet interfaces must be configured with either all FEC74 or all FEC91-KR4 mode.</p>

[Table 97 on page 273](#) summarizes the supported port speeds on MX10K-LC9600.

Table 97: Port Speed for MX10K-LC9600

PIC	Port Number	Port Speed Supported
PIC 0 through PIC 5	0-3	<p>4x10-Gigabit Ethernet</p> <p>4x25-Gigabit Ethernet</p> <p>2x100-Gigabit Ethernet</p> <p>4x100-Gigabit Ethernet</p> <p>Default: All active ports operate in 400-Gigabit Ethernet mode.</p>

Table 98: Interface Naming Convention for MX10K-LC9600

PIC	Interface Type	Interfaces
PIC 0	4x10-Gigabit Ethernet Interface	xe-0/0/0:0 – xe-0/0/0:3

Table 98: Interface Naming Convention for MX10K-LC9600 (Continued)

PIC	Interface Type	Interfaces
	4x25-Gigabit Ethernet Interface	et-0/0/0:0 – et-0/0/0:3
	4x100-Gigabit Ethernet Interface	
	40-Gigabit Ethernet Interface	et-0/0/0
	100-Gigabit Ethernet Interface	
	400-Gigabit Ethernet Interface	
	8x50-Gigabit Ethernet Interface	et-0/0/0:0 – et-0/0/0:7
	2x100-Gigabit Ethernet Interface	et-0/0/0:0 – et-0/0/0:1

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" [on page 177](#). For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" [on page 178](#).

Port Speed on PTX Routers

IN THIS SECTION

- [Port Speed on PTX10K-LC1201 | 275](#)
- [Port Speed on PTX10001-36MR Router | 283](#)
- [Port Speed on PTX10K-LC1202-36MR | 293](#)
- [Port Speed on PTX10003 | 301](#)

In this topic we have information related to the port speed on a PTX routers and line cards, support for multiple port speed details, guidelines and how to configure the port speed.

Port Speed on PTX10K-LC1201

PTX10008 routers support 8 PTX10K-LC1201 line cards. Starting in Junos OS Evolved Release 20.1R2 and 20.2R1, we now support a new port profile configuration to configure port speeds on the PTX10K-LC1201 line card. You can configure the port speed on the PTX10K-LC1201 line card by using the port profile configuration commands in the [edit interfaces] hierarchy. To streamline the configuration, the new port profile configuration commands are migrated from the [edit chassis] hierarchy to the [edit interfaces] hierarchy for the PTX10K-LC1201 line card.

Starting in Junos OS Evolved Release 20.4R1, you can configure 10Gbps speed and select the wavelength with SFP+ optics by plugging in the QSA adapter on the QSFP/QSFP+ ports of PTX10K-LC1201 line card. You can now configure the 10Gbps speed by using the `set interfaces et-<fpc>/<pic>/<port> speed <10g>` or `set interfaces et-<fpc>/<pic>/<port> speed <10g> number-of-sub-ports (1)` commands. You can select the wavelength of the SFP+ ports by using the `set interfaces et-<fpc>/<pic>/<port>:channel number optics-options wavelength wavelength` command and view the output using the `show interfaces diagnostics optics name` command.

For information on the line card, see [PTX10K-LC1201-36D for PTX10008 Routers](#).

To view the supported transceivers, optical interfaces, and DAC cables on PTX10K-LC1201, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 99 on page 275](#) summarizes the line card details and their description.

Table 99: PTX10K-LC1201 Software Support and Description

Software Feature Support	Description
Forward Error Correction (FEC)	<ul style="list-style-type: none"> By default, 200-Gigabit Ethernet interfaces and 400-Gigabit Ethernet interfaces support KP4 FEC mode. You cannot disable FEC mode explicitly. FEC mode is based on the type of optics connected. With non juniper optics, you can enable FEC mode by manually configuring it. When you manually configure the FEC mode, all the interfaces on a port must be configured with same FEC mode. For example, for a 4x100-Gigabit Ethernet configuration mode, each of the four 100-Gigabit Ethernet interfaces must be configured with either all FEC74 or all FEC91-KR4 mode.
Maximum Transmission Unit (MTU) Size	Maximum MTU of size 16000 bytes for transit traffic. Ingress interface limits is 9000 bytes.

Table 99: PTX10K-LC1201 Software Support and Description (Continued)

Software Feature Support	Description
Port Speed Configuration	<ul style="list-style-type: none"> For Junos OS Evolved Releases 19.4R1S1 and 20.1R1, you can configure the port speed at the [edit chassis] hierarchy. For Junos OS Evolved Releases 20.1R2, 20.2R1, and later, you can configure the port speed at the [edit interfaces] hierarchy.

When you change the speed of the port, or change the number-of-sub-ports per port, or configure or remove the number-of-sub-ports, the interfaces will be deleted and re-created for that port.

Table 100: Port Speed for PTX10K-LC1201

PIC	Port Number	Port Speeds Supported
PIC 0	0-35	1x10-Gigabit Ethernet 4x10-Gigabit Ethernet 1x40-Gigabit Ethernet 4x25-Gigabit Ethernet 8x25-Gigabit Ethernet 2x50-Gigabit Ethernet 1x100-Gigabit Ethernet 2x100-Gigabit Ethernet 4x100-Gigabit Ethernet 1x400-Gigabit Ethernet Default: All the active ports operate in 400-Gigabit Ethernet mode.

You can configure the active physical ports in a PIC. By default, all the physical ports in an interface are active. Use the `unused` command to control the number of interfaces created on a physical port in a PIC. If you configure a port as `unused`, no interfaces (channelized or non-channelized) are created for that port.

On channelized interfaces, the prefix `et-` is used irrespective of the speed you configure. All channelized interfaces have the same speed. You cannot configure an individual speed for each channelized

interface. You can configure the number of channelized interfaces by using the number-of-sub-ports command. On non-channelized interfaces, the prefix **et-** is used irrespective of the speed configured. If you do not configure the speed, by using the optional speed command, default speed is assigned to the interface.

Table 101: Interface Naming Convention for PTX10K-LC1201 line card (PTX10008 router)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	200-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
PIC 0	et-x/0/0: [0-7]	et-x/0/0: [0-7]	et-x/0/0	et-x/0/0: [0-7]	et-x/0/0 et-x/0/0:[0-1] et-x/0/0:[0-3]	et-x/0/0: [0-2]	et-x/0/0
	et-x/0/1: [0-7]	et-x/0/1: [0-7]	et-x/0/1	et-x/0/1: [0-7]	et-x/0/1 et-x/0/1:[0-1] et-x/0/1:[0-3]	et-x/0/1: [0-2]	et-x/0/1
	et-x/0/2: [0-7]	et-x/0/2: [0-7]	et-x/0/2	et-x/0/2: [0-7]	et-x/0/2 et-x/0/2:[0-1] et-x/0/2:[0-3]	et-x/0/2: [0-2]	et-x/0/2
	et-x/0/3: [0-7]	et-x/0/3: [0-7]	et-x/0/3	et-x/0/3: [0-7]	et-x/0/3 et-x/0/3:[0-1] et-x/0/3:[0-3]	et-x/0/3: [0-2]	et-x/0/3
	et-x/0/4: [0-7]	et-x/0/4: [0-7]	et-x/0/4	et-x/0/4: [0-7]	et-x/0/4 et-x/0/4:[0-1] et-x/0/4:[0-3]	et-x/0/4: [0-2]	et-x/0/4

Table 101: Interface Naming Convention for PTX10K-LC1201 line card (PTX10008 router) (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	200-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	et-x/0/5: [0-7]	et-x/0/5: [0-7]	et-x/0/5	et-x/0/5: [0-7]	et-x/0/5 et-x/0/5:[0-1] et-x/0/5:[0-3]	et-x/0/5: [0-2]	et-x/0/5
	et-x/0/6: [0-7]	et-x/0/6: [0-7]	et-x/0/6	et-x/0/6: [0-7]	et-x/0/6 et-x/0/6:[0-1] et-x/0/6:[0-3]	et-x/0/6: [0-2]	et-x/0/6
	et-x/0/7: [0-7]	et-x/0/7: [0-7]	et-x/0/7	et-x/0/7: [0-7]	et-x/0/7 et-x/0/7:[0-1] et-x/0/7:[0-3]	et-x/0/7: [0-2]	et-x/0/7
	et-x/0/8: [0-7]	et-x/0/8: [0-7]	et-x/0/8	et-x/0/8: [0-7]	et-x/0/8 et-x/0/8:[0-1] et-x/0/8:[0-3]	et-x/0/8: [0-2]	et-x/0/8
	et-x/0/9: [0-7]	et-x/0/9: [0-7]	et-x/0/9	et-x/0/9: [0-7]	et-x/0/9 et-x/0/9:[0-1] et-x/0/9:[0-3]	et-x/0/9: [0-2]	et-x/0/9
	et-x/0/10: [0-7]	et-x/0/10: [0-7]	et-x/0/10	et-x/0/10: [0-7]	et-x/0/10 et-x/0/10:[0-1] et-x/0/10:[0-3]	et-x/0/10: [0-2]	et-x/0/10

Table 101: Interface Naming Convention for PTX10K-LC1201 line card (PTX10008 router) (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	200-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	et-x/0/11: [0-7]	et-x/0/11: [0-7]	et-x/0/11	et-x/0/11: [0-7]	et-x/0/11 et-x/0/11:[0-1] et-x/0/11:[0-3]	et-x/0/11: [0-2]	et-x/0/11
	et-x/0/12: [0-7]	et-x/0/12: [0-7]	et-x/0/12	et-x/0/12: [0-7]	et-x/0/12 et-x/0/12:[0-1] et-x/0/12:[0-3]	et-x/0/12: [0-2]	et-x/0/12
	et-x/0/13: [0-7]	et-x/0/13: [0-7]	et-x/0/13	et-x/0/13: [0-7]	et-x/0/13 et-x/0/13:[0-1] et-x/0/13:[0-3]	et-x/0/13: [0-2]	et-x/0/13
	et-x/0/14: [0-7]	et-x/0/14: [0-7]	et-x/0/14	et-x/0/14: [0-7]	et-x/0/14 et-x/0/14:[0-1] et-x/0/14:[0-3]	et-x/0/14: [0-2]	et-x/0/14
	et-x/0/15: [0-7]	et-x/0/15: [0-7]	et-x/0/15	et-x/0/15: [0-7]	et-x/0/15 et-x/0/15:[0-1] et-x/0/15:[0-3]	et-x/0/15: [0-2]	et-x/0/15
	et-x/0/16: [0-7]	et-x/0/16: [0-7]	et-x/0/16	et-x/0/16: [0-7]	et-x/0/16 et-x/0/16:[0-1] et-x/0/16:[0-3]	et-x/0/16: [0-2]	et-x/0/16

Table 101: Interface Naming Convention for PTX10K-LC1201 line card (PTX10008 router) (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	200-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	et-x/0/17: [0-7]	et-x/0/17: [0-7]	et-x/0/17	et-x/0/17: [0-7]	et-x/0/17 et-x/0/17:[0-1] et-x/0/17:[0-3]	et-x/0/17: [0-2]	et-x/0/17
	et-x/0/18: [0-7]	et-x/0/18: [0-7]	et-x/0/18	et-x/0/18: [0-7]	et-x/0/18 et-x/0/18:[0-1] et-x/0/18:[0-3]	et-x/0/18: [0-2]	et-x/0/18
	et-x/0/19: [0-7]	et-x/0/19: [0-7]	et-x/0/19	et-x/0/19: [0-7]	et-x/0/19 et-x/0/19:[0-1] et-x/0/19:[0-3]	et-x/0/19: [0-2]	et-x/0/19
	et-x/0/20: [0-7]	et-x/0/20: [0-7]	et-x/0/20	et-x/0/20: [0-7]	et-x/0/20 et-x/0/20:[0-1] et-x/0/20:[0-3]	et-x/0/20: [0-2]	et-x/0/20
	et-x/0/21: [0-7]	et-x/0/21: [0-7]	et-x/0/21	et-x/0/21: [0-7]	et-x/0/21 et-x/0/21:[0-1] et-x/0/21:[0-3]	et-x/0/21: [0-2]	et-x/0/21
	et-x/0/22: [0-7]	et-x/0/22: [0-7]	et-x/0/22	et-x/0/22: [0-7]	et-x/0/22 et-x/0/22:[0-1] et-x/0/22:[0-3]	et-x/0/22: [0-2]	et-x/0/22

Table 101: Interface Naming Convention for PTX10K-LC1201 line card (PTX10008 router) (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	200-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	et-x/0/23: [0-7]	et-x/0/23: [0-7]	et-x/0/23	et-x/0/23: [0-7]	et-x/0/23 et-x/0/23:[0-1] et-x/0/23:[0-3]	et-x/0/23: [0-2]	et-x/0/23
	et-x/0/24: [0-7]	et-x/0/24: [0-7]	et-x/0/24	et-x/0/24: [0-7]	et-x/0/24 et-x/0/24:[0-1] et-x/0/24:[0-3]	et-x/0/24: [0-2]	et-x/0/24
	et-x/0/25: [0-7]	et-x/0/25: [0-7]	et-x/0/25	et-x/0/25: [0-7]	et-x/0/25 et-x/0/25:[0-1] et-x/0/25:[0-3]	et-x/0/25: [0-2]	et-x/0/25
	et-x/0/26: [0-7]	et-x/0/26: [0-7]	et-x/0/26	et-x/0/26: [0-7]	et-x/0/26 et-x/0/26:[0-1] et-x/0/26:[0-3]	et-x/0/26: [0-2]	et-x/0/26
	et-x/0/27: [0-7]	et-x/0/27: [0-7]	et-x/0/27	et-x/0/27: [0-7]	et-x/0/27 et-x/0/27:[0-1] et-x/0/27:[0-3]	et-x/0/27: [0-2]	et-x/0/27
	et-x/0/28: [0-7]	et-x/0/28: [0-7]	et-x/0/28	et-x/0/28: [0-7]	et-x/0/28 et-x/0/28:[0-1] et-x/0/28:[0-3]	et-x/0/28: [0-2]	et-x/0/28

Table 101: Interface Naming Convention for PTX10K-LC1201 line card (PTX10008 router) (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	200-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	et-x/0/29: [0-7]	et-x/0/29: [0-7]	et-x/0/29	et-x/0/29: [0-7]	et-x/0/29 et-x/0/29:[0-1] et-x/0/29:[0-3]	et-x/0/29: [0-2]	et-x/0/29
	et-x/0/30: [0-7]	et-x/0/30: [0-7]	et-x/0/30	et-x/0/30: [0-7]	et-x/0/30 et-x/0/30:[0-1] et-x/0/30:[0-3]	et-x/0/30: [0-2]	et-x/0/30
	et-x/0/31: [0-7]	et-x/0/31: [0-7]	et-x/0/31	et-x/0/31: [0-7]	et-x/0/31 et-x/0/31:[0-1] et-x/0/31:[0-3]	et-x/0/31: [0-2]	et-x/0/31
	et-x/0/32: [0-7]	et-x/0/32: [0-7]	et-x/0/32	et-x/0/32: [0-7]	et-x/0/32 et-x/0/32:[0-1] et-x/0/32:[0-3]	et-x/0/32: [0-2]	et-x/0/32
	et-x/0/33: [0-7]	et-x/0/33: [0-7]	et-x/0/33	et-x/0/33: [0-7]	et-x/0/33 et-x/0/33:[0-1] et-x/0/33:[0-3]	et-x/0/33: [0-2]	et-x/0/33
	et-x/0/34: [0-7]	et-x/0/34: [0-7]	et-x/0/34	et-x/0/34: [0-7]	et-x/0/34 et-x/0/34:[0-1] et-x/0/34:[0-3]	et-x/0/34: [0-2]	et-x/0/34

Table 101: Interface Naming Convention for PTX10K-LC1201 line card (PTX10008 router) (Continued)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	200-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface
	et-x/0/35: [0-7]	et-x/0/35: [0-7]	et-x/0/35	et-x/0/35: [0-7]	et-x/0/35 et-x/0/35:[0-1] et-x/0/35:[0-3]	et-x/0/35: [0-2]	et-x/0/35

For information about how to configure the speed of a port for a non-channelized interface using the new interfaces hierarchy, see "[Configure Port Speed for Non-Channelized Interfaces](#)" on page 179. For information about how to configure the speed of a port for a channelized interface using the new interfaces hierarchy, see "[Configure Port Speed for Channelized Interfaces](#)" on page 180.

Port Speed on PTX10001-36MR Router

For information on PTX10001-36MR hardware description, see *PTX10001-36MR Packet Transport Router Hardware Guide*.

To view the supported transceivers, optical interfaces, and DAC cables on PTX10001-36MR, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 102 on page 283](#) provides the basic details of PTX10001-36MR router.

Table 102: PTX10001-36MR Router Details and Description

Details	Description
PIC Level or Port Level Configuration	Only port level speed configuration is supported.
Port speed configuration:	Supports only [edit interfaces interface-name] hierarchy from Junos OS Evolved Release 20.3R1 onwards. Supports [edit chassis] hierarchy in Junos OS Evolved 20.2 releases.

Table 102: PTX10001-36MR Router Details and Description (Continued)

Details	Description
MTU size	<p>PTX10001-36MR router WAN interfaces support maximum MTU of size 16000 bytes for transit traffic. However for the traffic that is destined to host or is originating from host (such as protocol traffic), the maximum MTU limit is 9500 bytes. If any of the host bound packets or host originating packet is above 9500 bytes, then the packet will be dropped. Hence, if you are expecting host packets greater than 9500 bytes, then the WAN interface MTU value must be set as value less than or equal to 9500 bytes.</p>
Forward Error Correction (FEC) support	<ul style="list-style-type: none"> By default, KP4 FEC is enabled on 400 and 200-Gigabit Ethernet interfaces. Since it is mandatory to enable KP4 FEC option for 200 and 400-Gigabit Ethernet interfaces, you cannot disable KP4 FEC explicitly. The FEC mode is based on the type of optics connected, but it in some cases (with non juniper optics) can be enabled through manual configuration as well. In such case, all the interfaces on a port must be configured with same FEC mode. For example, for a 4x100-Gigabit Ethernet configuration mode, each of the four 100-Gigabit Ethernet interfaces must be configured with either all FEC74 or all FEC91-KR4 mode. <p>For information about FEC support, see <i>fec (gigether)</i>.</p>

Table 103: Speed Capability of Ports

Port numbers (PIC/Port Format)	Speed capability
0/0, 0/1, 0/2, 0/3, 0/8, 0/9, 0/10, 0/11	400-Gbps
1/0, 1/1, 1/2, 1/3, 1/8, 1/9, 1/10, 1/11	
2/0, 2/1, 2/2, 2/3, 2/8, 2/9, 2/10, 2/11	
0/4, 0/5, 0/6, 0/7	100-Gbps
1/4, 1/5, 1/6, 1/7	
2/4, 2/5, 2/6, 2/7	

If the number-of-sub-ports, do not match the type of optic connected and the configuration is invalid, then even though the interfaces are created, the links would remain down. A syslog entry will be added indicating usage of the wrong optic type.

Table 104: Channelization and FEC support on 400- and 200-Gigabit Ethernet capable ports

Speed Supported	You can channelize to:
400-Gigabit Ethernet capable ports	<ul style="list-style-type: none"> One or two 200-Gigabit Ethernet interfaces (1x200G or 2x200G) Four 100-Gigabit Ethernet interfaces or two 100-Gigabit Ethernet interfaces (4x100G or 2x100G) Two 50-Gigabit Ethernet interfaces or eight 50-Gigabit Ethernet interfaces (2x50G or 8x50G) Four 25-Gigabit Ethernet interfaces or eight 25-Gigabit Ethernet interfaces (4x25G or 8x25G) One 40-Gigabit Ethernet interface (1x40G) Four 10-Gigabit Ethernet interfaces (4x10G)
100-Gigabit Ethernet capable ports	<p>You can configure all four 100-Gigabit Ethernet capable ports (ports 4, 5, 6, and 7) to operate at 100-Gbps speeds.</p> <p>Note that, you can configure port 4 and port 6 to operate at 1x100-Gbps, 4x10-Gbps, 4x25-Gbps, or 1x40-Gbps speeds with the following conditions:</p> <ul style="list-style-type: none"> When port 4 is configured in 4x10-Gbps, 4x25-Gbps, or 1x40-Gbps, then port 5 must be configured as 'unused'. When port 6 is configured in 4x10-Gbps, 4x25-Gbps, or 1x40-Gbps, then port 7 must be configured as 'unused'.

Starting in Junos OS Evolved Release 20.3R1, you can configure the port speed on the PTX10001-36MR router by using the port profile configuration commands in the [edit interfaces] hierarchy. To streamline the configuration, the port profile configuration commands are migrated from the [edit chassis] hierarchy to the [edit interfaces] hierarchy for the PTX10001-36MR router.

From Junos OS Evolved Release 21.2R1, we support autonegotiation and linktraining using 400G DAC cable for PTX10001-36MR. This is not a default behavior. Use auto-negotiation under [edit interfaces interface-name ether-options] hierarchy level to enable autonegotiation and link training. Enabling autonegotiation automatically enables link training. This advertises configured speed on the interface hierarchy. It does not upgrade or downgrade speed based on peers advertised capability.

Autonegotiation and link training is supported only on the following channelization:

- One 400-Gigabit Ethernet interfaces (1x400G)
- Two 200-Gigabit Ethernet interfaces (2x200G)
- Four 100-Gigabit Ethernet interfaces (4x100G)

Configure speed at Port Level

To configure the PTX10001-36MR router at port level, follow the configuration steps in "[Configure Port Speed - Port Level](#)" on page 178 in 20.2 releases. See *speed* for more details.

From Junos OS Evolved Release 20.3R1 you must use, `set interfaces et-<fpc>/<pic>/<port> speed <10g | 25g | 50g | 40g | 100g | 200g | 400g>`, or `set interfaces et-<fpc>/<pic>/<port> speed <10g | 25g | 50g | 40g | 100g | 200g | 400g> number-of-sub-ports (1/2/4/8)` commands to configure the speed on the ports.

You can configure port profiles in the command line interface without the physical presence of an FPC. If an invalid port profile configuration is detected while booting a FPC, an alarm is generated. Also, the default port profile is selected for that PIC.

If the port profile configuration is changed while the FPC is up and running, and the new configuration is invalid, an alarm is generated. The existing port profile configured continues to be used for that PIC.

You can configure any supported speed on 400-Gbps capable ports. Configuring speed on one of the 400-Gbps capable port will not disrupt the traffic on any other ports.

But, for 100-Gbps capable ports only speeds of 100-Gbps, 25-Gbps, 40-Gbps, and 10-Gbps are valid.

From Junos OS Evolved Release 20.3R1 you must use, the following command:

```
set interfaces et-<fpc>/<pic>/<port> unused
```

If you configure *request security system-keystore unlink* command for a physical port in a PIC, then no channelized and non-channelized interfaces will be created. See *request security system-keystore unlink* for more information.

The [Table 107 on page 288](#) specifies which ports must be marked unused. If this rule is violated, then an alarm indicating port speed configuration error is raised. In such case, the existing running configuration will continue to be applied on such ports.

If the router reboots with such an invalid configuration, then the port with 40-Gbps, 4x10-Gbps, or 4x25-Gbps speed configuration and its counterpart port will not have any interfaces created for them.

[Table 105 on page 287](#) provides you the ports that you must power-off while configuring different speeds.

Table 105: Unused Port Settings

Port Speed	Valid ports to set the Port Speed	Ports you must explicitly power off (FPC/PIC/Port)
100-Gbps	All 100G capable ports can operate that 100-Gbps speed. You are not required to power off any 100G cable ports, in this case.	NA
40-Gbps, 25-Gbps, and 10-Gbps	0/0/4 and 0/0/6	0/0/5 and 0/0/7
	0/1/4 and 0/1/6	0/1/5 and 0/1/7
	0/2/4 and 0/2/6	0/2/5 and 0/2/7



NOTE: You can configure a port with more than one type of channelization mode for a given speed. For example, you can channelize a port to 1x100-Gbps, 2x100-Gbps, or 4x100-Gbps to configure the port to operate in 100-Gbps speed.

From Junos OS Evolved Release 20.3R1 to specify which of these channelization modes on the ports, you must execute the following command at the [edit interfaces] hierarchy:

```
set interfaces et-<fpc>/<pic>/<port> speed <10g | 25g | 50g | 40g | 100g | 200g | 400g> number-of-sub-ports (1/2/4/8)
```

See *number-of-sub-ports (interface)* for information on how to operate at different channelized speed.

When the number-of-sub-ports are not specified, the number of channels are created as per [Table 106 on page 287](#):

Table 106: Number of sub-ports supported for a particular speed

Port speed	Valid values for number-of-sub-ports
200-Gbps	1, 2
100-Gbps	1, 2, 4
QSFP56-DD-400G-ZR-M and QSFP56-DD-400G-ZR-M-HP support 3x100G channelisation.	3

Table 106: Number of sub-ports supported for a particular speed (Continued)

Port speed	Valid values for number-of-sub-ports
40-Gbps	1
50-Gbps	2, 8
25-Gbps	4, 8
10-Gbps	4

- The router has a fabric capacity of 9.6 Tbps. When the router operates in 24 400 Gigabit Ethernet mode, the fabric does not experience an oversubscription. The oversubscription occurs when the router is configured in 24 400-Gigabit Ethernet and 12 100-Gigabit Ethernet modes. To avoid oversubscription of a PIC, the sum of the total speed configured on ports 4 to 11 within the PIC must be less than 1.6 Tbps.

Table 107: Interface Naming Convention for PTX10001-36MR Router

PIC	4x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface	1x400-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	1x40-Gigabit Ethernet Interface
0	et-0/0/0: [0-3]	et-0/0/0: [0-3]	et-0/0/0: [0-1]	et-0/0/0: [0-7]	et-0/0/0	et-0/0/0	et-0/0/0
	et-0/0/1: [0-3]	et-0/0/1: [0-3]	et-0/0/1: [0-1]	et-0/0/1: [0-7]	et-0/0/1	et-0/0/1	et-0/0/1

Table 107: Interface Naming Convention for PTX10001-36MR Router (*Continued*)

PIC	4x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface	1x400-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	1x40-Gigabit Ethernet Interface
	et-0/0/2: [0-3]	et-0/0/2: [0-3]	et-0/0/2: [0-1]	et-0/0/2: [0-7]	et-0/0/2	et-0/0/2	et-0/0/2
	et-0/0/3: [0-3]	et-0/0/3: [0-3]	et-0/0/3: [0-1]	et-0/0/3: [0-7]	et-0/0/3	et-0/0/3	et-0/0/3
	Unsupported	et-0/0/4: [0-3]	Unsupported	Unsupported	Unsupported	et-0/0/4	et-0/0/4
	Unsupported	Unused	Unsupported	Unsupported	Unsupported	et-0/0/5	Unused
	Unsupported	et-0/0/6: [0-3]	Unsupported	Unsupported	Unsupported	et-0/0/6	et-0/0/6
	Unsupported	Unused	Unsupported	Unsupported	Unsupported	et-0/0/7	Unused
	et-0/0/8: [0-3]	et-0/0/8: [0-3]	et-0/0/8: [0-1]	et-0/0/8: [0-7]	et-0/0/8	et-0/0/8	et-0/0/8
	et-0/0/9: [0-3]	et-0/0/9: [0-3]	et-0/0/9: [0-1]	et-0/0/9: [0-7]	et-0/0/9	et-0/0/9	et-0/0/9

Table 107: Interface Naming Convention for PTX10001-36MR Router (Continued)

PIC	4x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface	1x400-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	1x40-Gigabit Ethernet Interface
	et-0/0/10: [0-3]	et-0/0/10: [0-3]	et-0/0/10: [0-1]	et-0/0/10: [0-7]	et-0/0/10	et-0/0/10	et-0/0/10
	et-0/0/11: [0-3]	et-0/0/11: [0-3]	et-0/0/11: [0-1]	et-0/0/11: [0-7]	et-0/0/11	et-0/0/11	et-0/0/11
1	et-0/1/0: [0-3]	et-0/1/0: [0-3]	et-0/1/0: [0-1]	et-0/1/0: [0-7]	et-0/1/0	et-0/1/0	et-0/1/0
	et-0/1/1: [0-3]	et-0/1/1: [0-3]	et-0/1/1: [0-1]	et-0/1/1: [0-7]	et-0/1/1	et-0/1/1	et-0/1/1
	et-0/1/2: [0-3]	et-0/1/2: [0-3]	et-0/1/2: [0-1]	et-0/1/2: [0-7]	et-0/1/2	et-0/1/2	et-0/1/2
	et-0/1/3: [0-3]	et-0/1/3: [0-3]	et-0/1/3: [0-1]	et-0/1/3: [0-7]	et-0/1/3	et-0/1/3	et-0/1/3
	Unsupported	et-0/1/4: [0-3]	Unsupported	Unsupported	Unsupported	et-0/1/4	et-0/1/4
	Unsupported	Unused	Unsupported	Unsupported	Unsupported	et-0/1/5	Unused

Table 107: Interface Naming Convention for PTX10001-36MR Router (*Continued*)

PIC	4x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface	1x400-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	1x40-Gigabit Ethernet Interface
	Unsupported	et-0/1/6: [0-3]	Unsupported	Unsupported	Unsupported	et-0/1/6	et-0/1/6
	Unsupported	Unused	Unsupported	Unsupported	Unsupported	et-0/1/7	Unused
	et-0/1/8: [0-3]	et-0/1/8: [0-3]	et-0/1/8: [0-1]	et-0/1/8: [0-7]	et-0/1/8	et-0/1/8	et-0/1/8
	et-0/1/9: [0-3]	et-0/1/9: [0-3]	et-0/1/9: [0-1]	et-0/1/9: [0-7]	et-0/1/9	et-0/1/9	et-0/1/9
	et-0/1/10: [0-3]	et-0/1/10: [0-3]	et-0/1/10: [0-1]	et-0/1/10: [0-7]	et-0/1/10	et-0/1/10	et-0/1/10
	et-0/1/11: [0-3]	et-0/1/11: [0-3]	et-0/1/11: [0-1]	et-0/1/11: [0-7]	et-0/1/11	et-0/1/11	et-0/1/11
2	et-0/2/0: [0-3]	et-0/2/0: [0-3]	et-0/2/0: [0-1]	et-0/2/0: [0-7]	et-0/2/0	et-0/2/0	et-0/2/0
	et-0/2/1: [0-3]	et-0/2/1: [0-3]	et-0/2/1: [0-1]	et-0/2/1: [0-7]	et-0/2/1	et-0/2/1	et-0/2/1

Table 107: Interface Naming Convention for PTX10001-36MR Router (*Continued*)

PIC	4x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface	1x400-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	1x40-Gigabit Ethernet Interface
	et-0/2/2: [0-3]	et-0/2/2: [0-3]	et-0/2/2: [0-1]	et-0/2/2: [0-7]	et-0/2/2	et-0/2/2	et-0/2/2
	et-0/2/3: [0-3]	et-0/2/3: [0-3]	et-0/2/3: [0-1]	et-0/2/3: [0-7]	et-0/2/3	et-0/2/3	et-0/2/3
	Unsupported	et-0/2/4: [0-3]	Unsupported	Unsupported	Unsupported	et-0/2/4	et-0/2/4
	Unsupported	Unused	Unsupported	Unsupported	Unsupported	et-0/2/5	Unused
	Unsupported	et-0/2/6: [0-3]	Unsupported	Unsupported	Unsupported	et-0/2/6	et-0/2/6
	Unsupported	Unused	Unsupported	Unsupported	Unsupported	et-0/2/7	Unused
	et-0/2/8: [0-3]	et-0/2/8: [0-3]	et-0/2/8: [0-1]	et-0/2/8: [0-7]	et-0/2/8	et-0/2/8	et-0/2/8
	et-0/2/9: [0-3]	et-0/2/9: [0-3]	et-0/2/9: [0-1]	et-0/2/9: [0-7]	et-0/2/9	et-0/2/9	et-0/2/9

Table 107: Interface Naming Convention for PTX10001-36MR Router (Continued)

PIC	4x100-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface	1x400-Gigabit Ethernet Interface	1x100-Gigabit Ethernet Interface	1x40-Gigabit Ethernet Interface
			4x25-Gigabit Ethernet Interface	2x50-Gigabit Ethernet Interface	8x50-Gigabit Ethernet Interface	1x200-Gigabit Ethernet Interface	
				2x200-Gigabit Ethernet Interface			
	et-0/2/10: [0-3]	et-0/2/10: [0-3]	et-0/2/10: [0-1]	et-0/2/10: [0-7]	et-0/2/10	et-0/2/10	et-0/2/10
	et-0/2/11: [0-3]	et-0/2/11: [0-3]	et-0/2/11: [0-1]	et-0/2/11: [0-7]	et-0/2/11	et-0/2/11	

For information about how to configure the speed of a port for a non-channelized interface using the new interfaces hierarchy, see ["Configure Port Speed for Non-Channelized Interfaces" on page 179](#). For information about how to configure the speed of a port for a channelized interface using the new interfaces hierarchy, see ["Configure Port Speed for Channelized Interfaces" on page 180](#).

Port Speed on PTX10K-LC1202-36MR

PTX10K-LC1202-36MR is a new fixed-configuration line card with 36 built-in ports which you can install in PTX10008 routers. On the PTX10K-LC1202-36MR, you can choose to configure the line card:

- To operate at 3.6T mode, configure all the 36 ports to 100-Gbps speed, or
- To operate at 4.8T mode, configure the four 400 GE ports (4, 10, 24, and 30) to operate at 400-Gbps speed and the remaining ports to operate at 100-Gbps speed.

For information on the line card, see [PTX10008 Line Card Components and Descriptions](#).

To view the supported transceivers, optical interfaces, and DAC cables on PTX10K-LC1202-36MR, see [Hardware Compatibility Tool \(HCT\)](#).

Table 108: Port Speed for PTX10K-LC1202-36MR for PTX10008 Routers

PIC	Port Number	Port Speed Supported
PIC 0 (Logical PIC)	400G capable ports - 4, 10, 24, and 30	400-Gigabit Ethernet 4x100-Gigabit Ethernet 2x100-Gigabit Ethernet 40-Gigabit Ethernet 4x10-Gigabit Ethernet 8x25-Gigabit Ethernet 10-Gigabit Ethernet Default Speed: 400-Gigabit Ethernet
	100G capable ports - 0 - 35	100-Gigabit Ethernet 4x25-Gigabit Ethernet 4x10-Gigabit Ethernet 40-Gigabit Ethernet 10-Gigabit Ethernet Default Speed: 100-Gigabit Ethernet

"[Configure Port Speed for Non-Channelized Interfaces](#)" on page 179 and "[Configure Port Speed for Channelized Interfaces](#)" on page 180 describes the steps to configure the port speed for channelized and non-channelized interfaces from the [edit interfaces] hierarchy.

To configure ports at different speed, use the `set interfaces et-<fpc>/<pic>/<port> speed <10g | 25g | 40g | 100g | 400g> number-of-sub-ports <1|2|4|8>` command.

For the steps to configure the port speed from the [edit interfaces] hierarchy. See `speed` for more details.

To channelize an interface, use the `number-of-sub-ports` command. The `number-of-sub-ports` specifies the number of channelized interfaces that you can configure on a physical port with the specified speed. Each channel operates at the specified speed. The default value of `number-of-sub-ports` per optics is 1. See `number-of-sub-ports`.

When the `number-of-sub-ports` are not specified, the number of channels are created as per [Table 109 on page 295](#):

Table 109: Number of sub-ports supported for a particular speed

Port speed	Valid values for number-of-sub-ports
400-Gbps	1
100-Gbps	1, 2, 4
40-Gbps	1
25-Gbps	4, 8
10-Gbps	1, 4

You can configure channelization on port 0 to port 3 and port 18 to port 21 with the following guidelines:

- The 1x10-Gigabit Ethernet, 4x10-Gigabit Ethernet or 1x40-Gigabit Ethernet and 4x25-Gigabit Ethernet channelization is supported only on even numbered ports (that is, 0, 2, 18, 20).
- When the even port x is channelized, you must set the odd port x+1 as unused. You must explicitly configure the ports that must be set as unused, using the following command:

`set interface et-<fpc>/<pic>/<port> unused`

- You cannot configure 1x10-Gigabit Ethernet, 4x10-Gigabit Ethernet or 1x40-Gigabit Ethernet and 4x25-Gigabit Ethernet simultaneously on the following ports:
 - port 0 and 2
 - port 18 and 20

You can configure port profiles in the command line interface without the physical presence of an FPC. If an invalid port profile configuration is detected while booting a FPC, an alarm is generated. Also, the default port profile is selected for that PIC.

If the port profile configuration is changed while the FPC is up and running, and the new configuration is invalid, an alarm is generated. The existing port profile configured continues to be used for that PIC.

To configure FEC mode, see *fec (gigether)*.

You can configure every interface to loopback mode, see *loopback*.

Table 110 on page 296 lists the naming conventions used for interfaces on PTX10K-LC1202-36MR (for 100-Gbps, 2x100-Gbps, 400-Gbps, and 4x100-Gbps speeds) for PTX10008 routers. PTX10008 routers support eight PTX10K-LC1202-36MR line cards.

Table 110: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 100-Gbps, 2x100-Gbps, 400-Gbps, and 4x100-Gbps

PIC	100-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface	4x100-Gigabit Ethernet Interface
PIC 0	et-x/0/0	Not Supported	Not Supported	Not Supported
	et-x/0/1	Not Supported	Not Supported	Not Supported
	et-x/0/2	Not Supported	Not Supported	Not Supported
	et-x/0/3	Not Supported	Not Supported	Not Supported
	et-x/0/4	et-x/0/4:[0-1]	et-x/0/4	et-x/0/4:[0-3]
	et-x/0/5	Not Supported	Not Supported	Not Supported
	et-x/0/6	Not Supported	Not Supported	Not Supported
	et-x/0/7	Not Supported	Not Supported	Not Supported
	et-x/0/8	Not Supported	Not Supported	Not Supported
	et-x/0/9	Not Supported	Not Supported	Not Supported
	et-x/0/10	et-x/0/10:[0-1]	et-x/0/10	et-x/0/10:[0-3]
	et-x/0/11	Not Supported	Not Supported	Not Supported
	et-x/0/12	Not Supported	Not Supported	Not Supported

Table 110: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 100-Gbps, 2x100-Gbps, 400-Gbps, and 4x100-Gbps (Continued)

PIC	100-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface	4x100-Gigabit Ethernet Interface
	et-x/0/13	Not Supported	Not Supported	Not Supported
	et-x/0/14	Not Supported	Not Supported	Not Supported
	et-x/0/15	Not Supported	Not Supported	Not Supported
	et-x/0/16	Not Supported	Not Supported	Not Supported
	et-x/0/17	Not Supported	Not Supported	Not Supported
	et-x/0/18	Not Supported	Not Supported	Not Supported
	et-x/0/19	Not Supported	Not Supported	Not Supported
	et-x/0/20	Not Supported	Not Supported	Not Supported
	et-x/0/21	Not Supported	Not Supported	Not Supported
	et-x/0/22	Not Supported	Not Supported	Not Supported
	et-x/0/23	Not Supported	Not Supported	Not Supported
	et-x/0/24	et-x/0/24:[0-1]	et-x/0/24	et-x/0/24:[0-3]
	et-x/0/25	Not Supported	Not Supported	Not Supported
	et-x/0/26	Not Supported	Not Supported	Not Supported

Table 110: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 100-Gbps, 2x100-Gbps, 400-Gbps, and 4x100-Gbps (Continued)

PIC	100-Gigabit Ethernet Interface	2x100-Gigabit Ethernet Interface	400-Gigabit Ethernet Interface	4x100-Gigabit Ethernet Interface
	et-x/0/27	Not Supported	Not Supported	Not Supported
	et-x/0/28	Not Supported	Not Supported	Not Supported
	et-x/0/29	Not Supported	Not Supported	Not Supported
	et-x/0/30	et-x/0/30:[0-1]	et-x/0/30	et-x/0/30:[0-3]
	et-x/0/31	Not Supported	Not Supported	Not Supported
	et-x/0/32	Not Supported	Not Supported	Not Supported
	et-x/0/33	Not Supported	Not Supported	Not Supported
	et-x/0/34	Not Supported	Not Supported	Not Supported
	et-x/0/35	Not Supported	Not Supported	Not Supported

[Table 111 on page 298](#) lists the naming conventions used for interfaces on PTX10K-LC1202-36MR (for 10-Gbps, 4x10-Gbps, 40-Gbps, 4x25-Gbps, and 8x25-Gbps speeds) for PTX10008 routers. PTX10008 routers support eight PTX10K-LC1202-36MR line cards.

Table 111: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 10-Gbps, 4x10-Gbps, 4x25-Gbps, and 8x25-Gbps

PIC	10-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	4x25-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface
PIC 0	et-x/0/0	et-x/0/0:[0-3]	et-x/0/0	et-x/0/0:[0-3]	Not Supported

Table 111: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 10-Gbps, 4x10-Gbps, 4x25-Gbps, and 8x25-Gbps (Continued)

PIC	10-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	4x25-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et-x/0/2	et-x/0/2:[0-3]	et-x/0/2	et-x/0/2:[0-3]	Not Supported
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et-x/0/4	et-x/0/4:[0-3]	et-x/0/4	et-x/0/4:[0-3]	et-x/0/4:[0-7]
	et-x/0/5	et-x/0/5:[0-3]	et-x/0/5	et-x/0/5:[0-3]	Not Supported
	et-x/0/6	et-x/0/6:[0-3]	et-x/0/6	et-x/0/6:[0-3]	Not Supported
	et-x/0/7	et-x/0/7:[0-3]	et-x/0/7	et-x/0/7:[0-3]	Not Supported
	et-x/0/8	et-x/0/8:[0-3]	et-x/0/8	et-x/0/8:[0-3]	Not Supported
	et-x/0/9	et-x/0/9:[0-3]	et-x/0/9	et-x/0/9:[0-3]	Not Supported
	et-x/0/10	et-x/0/10:[0-3]	et-x/0/10	et-x/0/10:[0-3]	et-x/0/10:[0-7]
	et-x/0/11	et-x/0/11:[0-3]	et-x/0/11	et-x/0/11:[0-3]	Not Supported
	et-x/0/12	et-x/0/12:[0-3]	et-x/0/12	et-x/0/12:[0-3]	Not Supported
	et-x/0/13	et-x/0/13:[0-3]	et-x/0/13	et-x/0/13:[0-3]	Not Supported
	et-x/0/14	et-x/0/14:[0-3]	et-x/0/14	et-x/0/14:[0-3]	Not Supported

Table 111: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 10-Gbps, 4x10-Gbps, 4x25-Gbps, and 8x25-Gbps (Continued)

PIC	10-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	4x25-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface
	et-x/0/15	et-x/0/15:[0-3]	et-x/0/15	et-x/0/15:[0-3]	Not Supported
	et-x/0/16	et-x/0/16:[0-3]	et-x/0/16	et-x/0/16:[0-3]	Not Supported
	et-x/0/17	et-x/0/17:[0-3]	et-x/0/17	et-x/0/17:[0-3]	Not Supported
	et-x/0/18	et-x/0/18:[0-3]	et-x/0/18	et-x/0/18:[0-3]	Not Supported
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et-x/0/20	et-x/0/20:[0-3]	et-x/0/20	et-x/0/20:[0-3]	Not Supported
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et-x/0/22	et-x/0/22:[0-3]	et-x/0/22	et-x/0/22:[0-3]	Not Supported
	et-x/0/23	et-x/0/23:[0-3]	et-x/0/23	et-x/0/23:[0-3]	Not Supported
	et-x/0/24	et-x/0/24:[0-3]	et-x/0/24	et-x/0/24:[0-3]	et-x/0/24:[0-7]
	et-x/0/25	et-x/0/25:[0-3]	et-x/0/25	et-x/0/25:[0-3]	Not Supported
	et-x/0/26	et-x/0/26:[0-3]	et-x/0/26	et-x/0/26:[0-3]	Not Supported
	et-x/0/27	et-x/0/27:[0-3]	et-x/0/27	et-x/0/27:[0-3]	Not Supported
	et-x/0/28	et-x/0/28:[0-3]	et-x/0/28	et-x/0/28:[0-3]	Not Supported

Table 111: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 10-Gbps, 4x10-Gbps, 4x25-Gbps, and 8x25-Gbps (Continued)

PIC	10-Gigabit Ethernet Interface	4x10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	4x25-Gigabit Ethernet Interface	8x25-Gigabit Ethernet Interface
	et-x/0/29	et-x/0/29:[0-3]	et-x/0/29	et-x/0/29:[0-3]	Not Supported
	et-x/0/30	et-x/0/30:[0-3]	et-x/0/30	et-x/0/30:[0-3]	et-x/0/30:[0-7]
	et-x/0/31	et-x/0/31:[0-3]	et-x/0/31	et-x/0/31:[0-3]	Not Supported
	et-x/0/32	et-x/0/32:[0-3]	et-x/0/32	et-x/0/32:[0-3]	Not Supported
	et-x/0/33	et-x/0/33:[0-3]	et-x/0/33	et-x/0/33:[0-3]	Not Supported
	et-x/0/34	et-x/0/34:[0-3]	et-x/0/34	et-x/0/34:[0-3]	Not Supported
	et-x/0/35	et-x/0/35:[0-3]	et-x/0/35	et-x/0/35:[0-3]	Not Supported

Port Speed on PTX10003

Starting in Junos OS Evolved Release 19.3R1, you can configure 400-gigabit ethernet interfaces using QSFP56-DD-400GBASE-LR8 optics on PTX10003 routers. Only ports 0, 4, 5, 9 within each logical PIC support 400-Gigabit ethernet mode. When using 400G on port 0, the total bandwidth (speed x number-of-subports) of port 1 has to be less than 100G and port 2 has to be configured as 'unused' (see *request security system-keystore unlink* for more details). When using port 4 as 400G, port 3 must be configured with total bandwidth of less than 100G and port 2 must be configured 'unused'. Similarly, with port 5, 9 using 400G, port 6, 8 respectively must be configured for less than 100G and port 7 should be configured as 'unused'. That is, when a port is configured in 400-Gigabit ethernet mode, you cannot configure speed of the adjacent port to be more than 100-Gbps, and the middle port (2 between 0~4 or 7 between 5~9) must be set to unused. For example, you can set et-0/0/0 to 400G, et-0/0/1 to 100G or less, but et-0/0/2 must be set to unused.

For information on the PTX10003 router, see [PTX10003-80C and PTX10003-160C Fixed Packet Transport Router Hardware Guide](#).

To view the supported transceivers, optical interfaces, and DAC cables on PTX10003, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 112 on page 302](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds for PTX10003-80C.

Table 112: Port Speed for PTX10003-80C

FPC	PIC	Ports	Port Speed Supported
FPC0	PIC0 PIC1	0-9 each (Port groups of 5 ports)	10-Gigabit Ethernet 25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 200-Gigabit Ethernet 400-Gigabit Ethernet Default: Dependent on the transceiver
FPC1	PIC0 PIC1	0-9 each (Port groups of 5 ports)	10-Gigabit Ethernet 25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 200-Gigabit Ethernet 400-Gigabit Ethernet Default: Dependent on the transceiver

When port speed and sub-port-number are configured, the configured values override the default port speed for the transceiver. If you try to configure a port speed that is not supported by the transceiver, the port will be disabled. If the port speed is not configured on a valid optical port, the PTX10003 router uses a default port speed of 2x100 Gbps. If the number-of-subports is not configured, 1x 40G | 100G |200G |400G data rate is assumed. 1x10G sub-port is not supported.

Table 113: Port Speed for PTX10003-160C

FPC	PIC	Ports	Port Speed Supported
FPC0	PIC0 PIC1	0-9 each (Port groups of 5 ports)	10-Gigabit Ethernet 25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 200-Gigabit Ethernet 400-Gigabit Ethernet Default: Dependent on the transceiver
FPC1	PIC0 PIC1	0-9 each (Port groups of 5 ports)	10-Gigabit Ethernet 25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 200-Gigabit Ethernet 400-Gigabit Ethernet Default: Dependent on the transceiver
FPC2	PIC0 PIC1	0-9 each (Port groups of 5 ports)	10-Gigabit Ethernet 25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 200-Gigabit Ethernet 400-Gigabit Ethernet Default: Dependent on the transceiver

Table 113: Port Speed for PTX10003-160C (Continued)

FPC	PIC	Ports	Port Speed Supported
FPC3	PIC0	0-9 each	10-Gigabit Ethernet
	PIC1	(Port groups of 5 ports)	25-Gigabit Ethernet 40-Gigabit Ethernet 100-Gigabit Ethernet 200-Gigabit Ethernet 400-Gigabit Ethernet Default: Dependent on the transceiver

The center port in each port group (port 2 and port 7) do not support 1x200 Gbps. To configure the speed as 200 Gbps for those ports, you'll need to configure them as 2x100 Gbps. For more details, see [Understanding QSFP-DD Interfaces and Configurations](#). Also, only ports 0, 4, 5, or 9 on each PIC can support 400 Gbps or 4x100 Gbps. To configure the speed, you must plug in the appropriate transceiver and configure the speed.

Table 114: Channelization Configuration on PTX10003

QSFP Transceiver	Native Port Speeds	Channelization Options
QSFP56-DD	1x400 Gbps	4x100 Gbps
QSFP28-DD	1x200 Gbps 2x100 Gbps	8x25 Gbps
QSFP28	1x100 Gbps	4x25 Gbps
QSFP+	1x40 Gbps 4x10 Gbps	4x10 Gbps

To configure the number of optical channels for a particular port if the optics are used in a channelized mode. You can use this configuration option to configure a speed (10, 40, and 100 Gbps) in different

number of channels based on the optics used. The default value of number-of-sub-ports per optics is 1. Following are the configurable values for the corresponding optic types:

Optic Type	QSFP56- DD-400GBASE- LR8 (400G)	QSFP DD 28F (200G)	QSFP 28 (100G)	QSFP+ (40G)	QSFP 28 (25G)	QSFP 28 DD(25G)	Default
Channelized	1	2	1	4	4	8	1
Non- channelized	NA	NA	1	1	1	1	1

You are not required to set any value for number-of-sub-ports while configuring 40G or 400G, as the default value for number-of-sub-ports is 1.

Table 115: Supported Port Speed and Unused Ports

Port Number	Speed
Port 0	400 Gbps
Port 1	100 Gbps
Port 2	Unused
Port 3	100 Gbps
Port 4	400 Gbps
Port 5	400 Gbps
Port 6	100 Gbps
Port 7	Unused
Port 8	100 Gbps

Table 115: Supported Port Speed and Unused Ports (*Continued*)

Port Number	Speed
Port 9	400 Gbps

Table 116: Configuration Guidelines for 4x100 Gbps (PTX10003)

Ports with speed of 4x100 Gbps	Ports configured as Unused
Port 0	Port 1 and Port 2
Port 4	Port 2 and Port 3
Port 5	Port 6 and Port 7
Port 9	Port 7 and Port 8

Table 117: Interface Naming Convention for PTX10003

PIC	4x10-Gigabit Ethernet Interface 4x25-Gigabit Ethernet interface	4x100-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	2x100 Gigabit Ethernet Interface	200-Gigabit Ethernet Interface 400-Gigabit Ethernet Interface
0	et-x/0/0 : [0-3]	et-x/0/0 : [0-3]	et-x/0/0	et-x/0/0	et-x/0/0 : [0-1]	et-x/0/0
	et-x/0/1 : [0-3]	et-x/0/1 : [0-3]	et-x/0/1	et-x/0/1	et-x/0/1 : [0-1]	et-x/0/1
	et-x/0/2 : [0-3]	et-x/0/2 : [0-3]	et-x/0/2	et-x/0/2	et-x/0/2 : [0-1]	et-x/0/2

Table 117: Interface Naming Convention for PTX10003 (Continued)

PIC	4x10-Gigabit Ethernet Interface	4x100-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface	2x100 Gigabit Ethernet Interface	200-Gigabit Ethernet Interface
	et-x/0/3 : [0-3]	et-x/0/3 : [0-3]	et-x/0/3	et-x/0/3	et-x/0/3 : [0-1]	et-x/0/3
	et-x/0/4 : [0-3]	et-x/0/4 : [0-3]	et-x/0/4	et-x/0/4	et-x/0/4 : [0-1]	et-x/0/4
	et-x/0/5 : [0-3]	et-x/0/5 : [0-3]	et-x/0/5	et-x/0/5	et-x/0/5 : [0-1]	et-x/0/5
	et-x/0/6 : [0-3]	et-x/0/6 : [0-3]	et-x/0/6	et-x/0/6	et-x/0/6 : [0-1]	et-x/0/6
	et-x/0/7 : [0-3]	et-x/0/7 : [0-3]	et-x/0/7	et-x/0/7	et-x/0/7 : [0-1]	et-x/0/7
	et-x/0/8 : [0-3]	et-x/0/8 : [0-3]	et-x/0/8	et-x/0/8	et-x/0/8 : [0-1]	et-x/0/8
	et-x/0/9 : [0-3]	et-x/0/9 : [0-3]	et-x/0/9	et-x/0/9	et-x/0/9 : [0-1]	et-x/0/9

For information on how to configure the speed at the PIC level, see "[Configure Port Speed - PIC Level](#)" on page 177. For information on how to configure the speed at the port level, see "[Configure Port Speed - Port Level](#)" on page 178.

5

PART

800ZR and 800G OpenZR+ Optical Transceivers

- Overview | [309](#)
- Features of 800ZR and 800G OpenZR+ | [310](#)
- Configure and Monitor 800ZR and 800G OpenZR+ Coherent Optics | [323](#)

Overview

SUMMARY

Gives an overview of 800ZR and 800G OpenZR+ optical transceivers. This topic also includes benefits of 800ZR and 800G OpenZR+ optics, and licensing information.

IN THIS SECTION

- Benefits of 800ZR and 800G OpenZR+ Optics | [310](#)
- Licensing | [310](#)

The 800ZR is an Optical Interworking Forum (OIF) standard for transporting 800GbE over long distances through fiber optic cables. The standard defines the solution to transport 800GbE optical signals using a single wavelength.

The primary application of 800ZR optics is to interconnect data centers within a distance of up to 450 km. Data centers require high bandwidth, high speed, and low latency networks. The 800ZR optical module meets these requirements.

800G OpenZR+ is an extension of 800ZR specifications to improve performance and adheres to OpenZR+ MSA standard.

Some characteristics of the 800ZR and 800G OpenZR+ optics are:

- Configuration of different speeds and negotiation between capacity and reach. For example, 800 Gbps can support nominal reach of 450 km, 400 Gbps up to 500 km, 300 Gbps up to 900 km, and 200 Gbps up to 2000 km.
- Capacity to carry the signals across long distances through fiber optic cables with periodic optical amplification.
- Fully C-band tunable and dense wavelength-division multiplexing (DWDM). We support 6.25/12.5/50/75/100 GHz channel spacings.
- Ability to operate in amplified and unamplified modes.

To know more about the Juniper Networks® 800G optical transceivers and cables, their specifications, and how to install, remove, and maintain these transceivers, see [Juniper 800G Optical Transceivers and Cables Guide](#).

Benefits of 800ZR and 800G OpenZR+ Optics

- **Cost efficiency:** Operators can use 800ZR and 800G OpenZR+ coherent DWDM pluggable optics in IP routers and switches, and directly interconnect to a DWDM line system. The use of these transceivers eliminates the need for external optical transponders, thereby reducing infrastructure maintenance costs and operating expenses for the operators.

- **Reduced carbon footprint, and power and space consumption:**

800ZR and 800G OpenZR+ optics reduce power consumption by up to 50 percent by eliminating the need for external transponders. The elimination of transponders also results in reduced carbon footprint and space requirement.

- **Increased system efficiency:**

800ZR and 800G OpenZR+ optics simplify operations with standard APIs and open data model.

Licensing

You need a license to use the 800ZR and 800G OpenZR+ standards. For information about the licensing model for the coherent 800ZR and 800G OpenZR+ optics, see [Software Licenses for Juniper Coherent Optics](#). For additional information about License management, see [Licensing Guide](#). You can also contact your Juniper Account Team or Juniper Partner.

Features of 800ZR and 800G OpenZR+

SUMMARY

Describes the 800ZR and 800G OpenZR+ features like application selection, transmit output power, loopback, and so on. This topic also includes the configuration steps of these features.

IN THIS SECTION

- Application Selection for 800G OpenZR + | 311
- Configure Application Selection on 800G OpenZR+ Optics | 315
- Transmit Optical Power for 800G OpenZR + | 319

- [Configure Transmit Output Power | 320](#)
- [Optical Loopback | 322](#)
- [Signal Power | 323](#)

Application Selection for 800G OpenZR+

SUMMARY

Describes about application selection on 800G OpenZR+ optics and its configuration.

The 800ZR and 800G OpenZR+ optics advertises supported applications for a particular speed. You can select any of the application based on your requirement. You can also switch between these applications. See [Table 118 on page 311](#) for devices that support application selection.

Table 118: Supported Applications of 800G OpenZR+ Optics

Application	Host Interface Code Description	Media Interface Code Description	Speed
1	800GAUI-8 S C2M (Annex 120G)(81 or 0x51)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	800G
2	800GAUI-8 L C2M (Annex 120G)(82 or 0x52)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	800G
3	400GAUI-4-S C2M (Annex 120G)(79 or 0x4F)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	2x400G

Table 118: Supported Applications of 800G OpenZR+ Optics (*Continued*)

Application	Host Interface Code Description	Media Interface Code Description	Speed
4	400GAUI-4-L C2M (Annex 120G)(80 or 0x50)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	2x400G
5	200GAUI-2-S C2M (Annex 120G)(77 or 0x4D)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	4x200G
6	200GAUI-2-L C2M (Annex 120G)(78 or 0x4E)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	4x200G
7	100GAUI-1-S C2M (Annex 120G)(75 0x4B)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	8x100G
8	100GAUI-1-L C2M (Annex 120G)(76 or 0x4C)	800ZR-A, 150 GHz DWDM (108 or 0x6C)	8x100G
9	200GAUI-2-L C2M (Annex 120G)(77 or 0x4D)	OpenZR+ ZR600-OFEC-8QAM(212)	3x200G
10	200GAUI-2-L C2M (Annex 120G)(78 0x4E)	OpenZR+ ZR600-OFEC-8QAM(212)	3x200G
11	100GAUI-1-S C2M (Annex 120G)(75 0x4B)	OpenZR+ ZR600-OFEC-8QAM(212)	6x100G
12	100GAUI-1-L C2M (Annex 120G)(76 0x4C)	OpenZR+ ZR600-OFEC-8QAM(212)	6x100G
13	400GAUI-8 C2M (Annex 120E)(17 or 0x11)	OpenZR+ ZR400-OFEC-QPSK(213)	1x400G
14	400GAUI-8 C2M (Annex 120E)(17 or 0x11)	ZR400-OFEC-8QAM-HA(55)	1x400G

Table 118: Supported Applications of 800G OpenZR+ Optics (Continued)

Application	Host Interface Code Description	Media Interface Code Description	Speed
15	200GAUI-4 C2M (Annex 120E)(15 or 0xF)	ZR400-OFEC-8QAM-HA(55)	2x200G

For example:

To display the configurable applications

- Use the `show interfaces diagnostics optics-applications <interface>` command to display the configurable applications.

```
user@host> show interfaces diagnostics optics-applications et-0/0/12:0

Physical interface: et-0/0/12:0

  Interface Name          : et-0/0/12:0
  Current Speed          : 2x400G

  Ap Sel     Host Intf Code          Media Intf
  Code          Host Lanes      Media Lanes      Host Assign
  Media Assign

  1      800GAUI-8 S C2M (Annex 120G)(81)      800ZR-A, 150 GHz
  DWDM(108)          8          1          1          1
  2      800GAUI-8 L C2M (Annex 120G)(82)      800ZR-A, 150 GHz
  DWDM(108)          8          1          1          1
  3      400GAUI-4-S C2M (Annex 120G)(79)      800ZR-A, 150 GHz
  DWDM(108)          4          1          17          1
  4      400GAUI-4-L C2M (Annex 120G)(80)      800ZR-A, 150 GHz
  DWDM(108)          4          1          17          1
  5      200GAUI-2-S C2M (Annex 120G)(77)      800ZR-A, 150 GHz
  DWDM(108)          2          1          85          1
```

6	200GAUI-2-L C2M (Annex 120G)(78) DWDM(108)	2	1	800ZR-A, 150 GHz 85	1
7	100GAUI-1-S C2M (Annex 120G)(75) DWDM(108)	1	1	800ZR-A, 150 GHz 255	1
8	100GAUI-1-L C2M (Annex 120G)(76) DWDM(108)	1	1	800ZR-A, 150 GHz 255	1
9	200GAUI-2-S C2M (Annex 120G)(77) OFEC-8QAM(212)	2	1	OpenZR+ ZR600- 21	1
10	200GAUI-2-L C2M (Annex 120G)(78) OFEC-8QAM(212)	2	1	OpenZR+ ZR600- 21	1
11	100GAUI-1-S C2M (Annex 120G)(75) OFEC-8QAM(212)	1	1	OpenZR+ ZR600- 63	1
12	100GAUI-1-L C2M (Annex 120G)(76) OFEC-8QAM(212)	1	1	OpenZR+ ZR600- 63	1
13	400GAUI-8 C2M (Annex 120E)(17) QPSK(213)	8	1	OpenZR+ ZR400-OFEC- 1	1
14	400GAUI-8 C2M (Annex 120E)(17) HA(55)	8	1	ZR400-OFEC-8QAM- 1	1
15	200GAUI-4 C2M (Annex 120E)(15) HA(55)	4	1	ZR400-OFEC-8QAM- 17	1

Follow the guidelines below while configuring speed and application selection:

- If configured speed and application selection do not match, you receive a The configuration is not supported due to speed mismatch alarm.
- You can switch between short and long application selection code for the same speed.
- If you configure Nx200G, then you receive a Invalid Port Speed Configuration, Refer 'show chassis pic' alarm. Configuration of host code, media code, and application selection id is not accepted and error message is logged.

- For 1x400G speed, application selection ID 14 with media code ZR400-OFEC-8QAM-HA is the default choice. If you want to configure 1x400G with media code OpenZR+ ZR400-OFEC-QPSK (213), you can do it by host/media/application selection code configuration.
- If the host side is configured with 8x100G and the line side with 1x800G, the link status mismatch is expected.

Configure Application Selection on 800G OpenZR+ Optics

SUMMARY

Describes the steps to configure application selection on 800G OpenZR+ optical transceivers.

Execute the following steps to configure the 800G OpenZR+ optics:

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

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

2. Include the speed statement at the [edit interfaces interface-name] hierarchy level to configure the speed.

```
[edit interfaces interface-name]
user@host# set speed (400g | 100g);
```

3. Configure speed and channelization.

```
[edit interfaces interface-name]
user@host# set number-of-sub-ports <number-of-sub-ports>
```

To configure a different application for a speed, see [Step 5](#).

4. Display all configurable applications in the configured speed. You can use this information while executing [Step 5](#).

```
user@host> show interfaces diagnostics optics-applications et-0/0/12:0

Physical interface: et-0/0/12:0

  Interface Name          : et-0/0/12:0
  Current Speed          : 2x400G

  Ap Sel    Host Intf Code          Media Intf
  Code          Host Lanes      Media Lanes  Host Assign
  Media Assign

  1      800GAUI-8 S C2M (Annex 120G)(81)      800ZR-A, 150 GHz
  DWDM(108)          8          1          1          1
  2      800GAUI-8 L C2M (Annex 120G)(82)      800ZR-A, 150 GHz
  DWDM(108)          8          1          1          1
  3      400GAUI-4-S C2M (Annex 120G)(79)      800ZR-A, 150 GHz
  DWDM(108)          4          1          17         1
  4      400GAUI-4-L C2M (Annex 120G)(80)      800ZR-A, 150 GHz
  DWDM(108)          4          1          17         1
  5      200GAUI-2-S C2M (Annex 120G)(77)      800ZR-A, 150 GHz
  DWDM(108)          2          1          85         1
  6      200GAUI-2-L C2M (Annex 120G)(78)      800ZR-A, 150 GHz
  DWDM(108)          2          1          85         1
  7      100GAUI-1-S C2M (Annex 120G)(75)      800ZR-A, 150 GHz
  DWDM(108)          1          1          255        1
  8      100GAUI-1-L C2M (Annex 120G)(76)      800ZR-A, 150 GHz
  DWDM(108)          1          1          255        1
  9      200GAUI-2-S C2M (Annex 120G)(77)      OpenZR+ ZR600-
  OFEC-8QAM(212)          2          1          21         1
```

10	200GAUI-2-L C2M (Annex 120G)(78) OFEC-8QAM(212)	2	1	OpenZR+ ZR600- 21	1
11	100GAUI-1-S C2M (Annex 120G)(75) OFEC-8QAM(212)	1	1	OpenZR+ ZR600- 63	1
12	100GAUI-1-L C2M (Annex 120G)(76) OFEC-8QAM(212)	1	1	OpenZR+ ZR600- 63	1
13	400GAUI-8 C2M (Annex 120E)(17) QPSK(213)	8	1	OpenZR+ ZR400-OFEC- 1	1
14	400GAUI-8 C2M (Annex 120E)(17) HA(55)	8	1	ZR400-OFEC-8QAM- 1	1
15	200GAUI-4 C2M (Annex 120E)(15) HA(55)	4	1	ZR400-OFEC-8QAM- 17	1

5. Change the application.

- Change to the [edit interfaces interface-name optics-options] hierarchy.

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

- Execute either of the following commands to change the application. Both the commands are mutually exclusive and you cannot use them together.

```
user@host# set interfaces <interface-name> optics-options application hostid <hostid>
mediaid <mediaid> [ domainid <domainid>]
```

```
user@host# set interfaces <interface-name> optics-options appselid <appselid>
```

- Display the output after you switch to a new application. The following output shows **media id** and **host id** as **108** and **79** respectively.

```
user@host> show interfaces et-0/0/12:0
```

```

Physical interface: et-0/0/12:1, Administratively down, Physical link is Down
  Interface index: 1256, SNMP ifIndex: 831
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Disabled, Media type: Fiber
  Wavelength      : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code      : 800ZR-A, 150 GHz DWDM
  Host Code       : 400GAUI-4-S C2M (Annex 120G)
  Device flags    : Present Running Down Admin-Disabled
  Interface flags: Hardware-Down Down Admin-Test SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 9c:5a:80:16:89:c9, Hardware address: 9c:5a:80:16:89:c9
  Last flapped    : 2025-07-08 16:40:15 PDT (04:28:58 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Active alarms   : LINK
  Active defects  : LINK

  PCS statistics          Seconds
    Bit errors            0
    Errorred blocks       0
  Ethernet FEC Mode     FEC119
    FEC Codeword size    544
    FEC Codeword rate     0.945
  Ethernet FEC statistics Errors
    FEC Corrected Errors 32949
    FEC Uncorrected Errors 0
    FEC Corrected Errors Rate 3
    FEC Uncorrected Errors Rate 0
    PRE FEC BER          0
  Optic FEC Mode        OFEC
  Optic FEC statistics:
    Corrected Errors      38439428692716
    Uncorrected Words      20
    Corrected Error rate   2422134408
    Uncorrected Error rate 0
    Uncorrected Error Ratio 0.00e+00
    Corrected Error Ratio ( 16134 seconds average) 2.61e-03
  PRBS Mode : Disabled
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring        : Disable

```

```

Logical interface et-0/0/12:1.16386 (Index 1036) (SNMP ifIndex 847)
  Flags: Device-Down SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol multiservice, MTU: Unlimited
  Flags: None

```

Transmit Optical Power for 800G OpenZR+

SUMMARY

Describes transmit optical power and its ranges.

IN THIS SECTION

● [Benefits](#) | [319](#)

Optical Transmitter Power Management allows you to adjust the transmit optical power within a fiber-optic network, ensuring optimal signal strength and reliable network performance. This feature leverages a precise command to control the transmit (Tx) power settings, accommodating various vendor-specific modules with configurable ranges and resolution. By utilizing this functionality, you can adjust the power output to avoid issues like signal degradation or over-amplification. You can also tailor the settings to the specific hardware in use for enhanced network efficiency and troubleshooting.

Benefits

- Ensures optimal signal strength by allowing precise adjustments to the transmit optical power, thereby maintaining reliable network performance.
- Prevents signal degradation and over-amplification by enabling fine-tuning of power output, which helps avoid common issues in fiber-optic networks.
- Enhances compatibility and performance by providing vendor-specific configurations, ensuring that power settings can be tailored to the specifications of different optical modules.
- Facilitates efficient network planning and troubleshooting by allowing users to customize power settings based on specific ranges and resolutions, leading to more effective management of network operations
- Improves overall network efficiency by enabling detailed adjustments to Tx power, which helps in achieving better signal quality and reduces the likelihood of network disruptions.

Configure Transmit Output Power

SUMMARY

Describes the steps to configure transmit power.

When you configure transmit power, it is important to keep the value within a range. Keeping the transmit power value within the range ensures signal stability and network performance. The transmit power range varies between different vendors.

Execute the following steps to configure 800G OpenZR+:

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

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

2. Configure output power.

```
user@host# set interface et-<interface> optics-option tx-power <tx power>
```

3. Verify the target output power using the following command.

```
show interfaces diagnostics optics <interface>
```

```
root@user# run show interfaces diagnostics optics et-0/0/10  
Physical interface: et-0/0/10  
Module temperature : 41 degrees C / 105 degrees F  
Module voltage : 3.18 V  
Wavelength channel number : 28  
Wavelength setpoint : 1550.12 nm  
Tx dither : Disabled  
Frequency error : 0.00 GHz  
Wavelength error : 0.0 nm  
TEC fault alarm : False  
Wavelength unlocked alarm : False
```

```
Tx tune alarm : False
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 temperature high alarm threshold : 76 degrees C / 168 degrees F
Module temperature low alarm threshold : -6 degrees C / 21 degrees F
Module temperature high warning threshold : 73 degrees C / 163 degrees F
Module temperature low warning threshold : -3 degrees C / 26 degrees F
Module voltage high alarm threshold : 3.63 V
Module voltage low alarm threshold : 3.0 V
Module voltage high warning threshold : 3.599 V
Module voltage low warning threshold : 3.049 V
Laser bias current high alarm threshold : 524.28 mA
Laser bias current low alarm threshold : 0.0 mA
Laser bias current high warning threshold : 524.28 mA
Laser bias current low warning threshold : 0.0 mA
Laser output power high alarm threshold : 0.199 mW / -7.00 dBm
Laser output power low alarm threshold : 0.015 mW / -18.01 dBm
Laser output power high warning threshold : 0.158 mW / -7.99 dBm
Laser output power low warning threshold : 0.019 mW / -16.98 dBm
Laser rx power high alarm threshold : 2.511 mW / 4.00 dBm
Laser rx power low alarm threshold : 0.004 mW / -23.01 dBm
Laser rx power high warning threshold : 1.995 mW / 3.00 dBm
Laser rx power low warning threshold : 0.009 mW / -20.00 dBm
Laser temperature high alarm threshold : 76 degrees C / 168 degrees F
Laser temperature low alarm threshold : -6 degrees C / 21 degrees F
Laser temperature high warning threshold : 73 degrees C / 163 degrees F
Laser temperature low warning threshold : -3 degrees C / 26 degrees F
Lane 0
Laser bias current : 234.8 mA
Laser output power : 0.063 mW / -10.00 dBm ----->tx-power
```

Optical Loopback

SUMMARY

Describes about the supported loopback types options in 800G OpenZR+.

IN THIS SECTION

- [Benefits of Optical Loopback | 322](#)

An optical loopback involves testing the transmit and receive capabilities of a fiber-optic interface. This is typically performed by connecting the transmit port to the receive port using a known good fiber cable, ensuring that both ports are operational and the fiber-optic connection is functioning as expected.

The 800G OpenZR+ supports the following loopback types:

- host-side-in: Loop host side input back to output
- host-side-out: Loop host side output back to input
- media-side-in: Loop media side input back to output
- media-side-out : Loop media side output back into input

Benefits of Optical Loopback

- Enhance network diagnostics by allowing comprehensive testing of signal paths within optics modules, ensuring all internal connections are functioning correctly.
- Maintain high signal integrity by enabling precise verification of both host-side and media-side signal loops, reducing the risk of undetected errors.
- Increase network reliability by facilitating regular maintenance checks and immediate detection of issues, preventing potential network failures.

Use the following command to configure loopback options:

```
user@host# set interfaces et-<fpc>/<pic>/<port> optics-options loopback loopbacktype ?
```

Possible completions:

host-side-in	Loop host side input back to output
host-side-out	Loop host side output back to input
media-side-in	Loop media side input back to output

media-side-out	Loop media side output back to input
----------------	--------------------------------------

Signal Power

SUMMARY

Describes the command to get the signal power performance.

The `show interface diagnostics optics et-<x/x/x>` command shows the Rx signal power in the mW/dBm format.

Configure and Monitor 800ZR and 800G OpenZR+ Coherent Optics

SUMMARY

Includes configuration and monitoring procedures of 800ZR and 800G OpenZR+ coherent optics.

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

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

2. Include the speed statement at the [edit interfaces interface-name] hierarchy level to configure the speed.

```
[edit interfaces interface-name]
user@host# set speed (800g | 400g);
```

3. Configure speed and channelization.

```
[edit interfaces interface-name]
user@host# set number-of-sub-ports <number-of-sub-ports>
```

4. Configure a different application for a speed.

- a. Change to the [edit interfaces interface-name optics-options].

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

- b. Execute either of the following commands to change the application. Both the commands are mutually exclusive and you cannot use them together.

```
user@host# set interfaces <interface-name> optics-options application hostid <hostid>
mediaid <mediaid> [domainid
<domainid>]
```

```
user@host# set interfaces <interface-name> optics-options appselid <appselid>
```

5. [Optional] Configure output power in the [edit interfaces interface-name optics-options] hierarchy level. If in channelized mode, you must use the first sub-port to configure output power.

```
user@host# set interface et-<interface-name> optics-option tx-power <tx power>
```

Example:

```
user@host# set interfaces et-0/0/0:0 optics-options tx-power -10
```

6. Specify the wavelength of optics in nm. If in channelized mode, you must use the first sub-port to configure wavelength.

```
[edit interfaces interface-name optics-options]
user@host# set wavelength nm
```

Example:

```
user@host# set interfaces et-0/0/0:0 optics-options wavelength 1552.52
```

7. To view the current and historical performance metrics, use the `show interfaces transport pm` command. See [show interfaces transport pm](#).

6

PART

400ZR and 400G OpenZR+ Optical Transceivers

- Overview | [327](#)
- Features of 400ZR and 400G OpenZR+ | [329](#)
- Configure and Monitor 400ZR and 400G OpenZR+ Coherent Optics | [347](#)

Overview

IN THIS SECTION

- Benefits of 400ZR and 400G OpenZR+ Optics | [328](#)
- Licensing | [328](#)

The 400ZR is an Optical Interworking Forum (OIF) standard for transporting 400GbE over long distances through fiber optic cables. The standard defines the solution to transport 400GbE optical signals using a single wavelength.

Use [400ZR and 400G OpenZR+](#) to confirm platform and release support.

The primary application of 400ZR optics (model numbers QDD-400-ZR and JCO400-QDD-ZR) is to interconnect data centers within a distance of up to 120 km. Data centers require high bandwidth, high speed, and low latency networks. The 400ZR optical module meets these requirements.

400G OpenZR+ (model numbers: QDD-400 ZR-M, JCO400-QDD-ZR-M, and QDD-400G-ZR-M-HP) is an extension of 400ZR specifications to improve performance and adheres to OpenZR+ MSA standard.

Some characteristics of the 400ZR and 400G OpenZR+ optics are:

- Configuration of different speeds and negotiation between capacity and reach. For example, 400 Gbps can support nominal reach of 500 km, 300 Gbps up to 900 km, and 200 Gbps up to 2000 km. Exact reach depends on an optical line system that is in use (400G OpenZR+ optics only).
- Capacity to transport 4x100 Gbps and 1x400 Gbps.
- Capacity to carry the signals across long distances through fiber optic cables with periodic optical amplification.
- Fully C-band tunable and dense wavelength-division multiplexing (DWDM). We support 75-GHz and 100-GHz channel spacings.
- Ability to operate in amplified and unamplified modes.
- Support for 400 Gbps speed. You can use the port as one 400GbE interface or channelize it to operate as four 100GbE interfaces.

For product specification and compatibility, see [Hardware Compatibility Tool](#).

Benefits of 400ZR and 400G OpenZR+ Optics

- **Cost efficiency:** Operators can use 400ZR and 400G OpenZR+ coherent DWDM pluggable optics in IP routers and switches, and directly interconnect to a DWDM line system. The use of these transceivers eliminates the need for external optical transponders, thereby reducing infrastructure maintenance costs and OpEx for the operators.

- **Reduced carbon footprint, and power and space consumption:**

400ZR and 400G OpenZR+ optics reduce power consumption by up to 50 percent by eliminating the need for external transponders. The elimination of transponders also results in reduced carbon footprint and space requirement.

- **Increased system efficiency:**

400ZR and 400G OpenZR+ optics simplify operations with standard APIs and open data model.

- **Channelization support:** 400ZR supports channelization of 1x400 Gbps and 4x100 Gbps speeds (muxponder mode). 400G OpenZR+ supports the following channelizations:

- 1x400 Gbps
- 4x100 Gbps
- 3x100 Gbps
- 2x100 Gbps
- 1x100 Gbps

As you decrease speed, your reach capability increases and power consumption decreases.

- **0dBm output transmit power:** QDD-400G-ZR-M-HP models can operate in high power mode and deliver a transmit power output of 0 dBm.

Licensing

You need a license to use the 400ZR and 400G OpenZR+ standards. For information about the licensing model for the coherent 400ZR and 400G OpenZR+ optics, see [Software Licenses for Juniper Coherent Optics](#). For additional information about License management, see [Licensing Guide](#). You can also contact your Juniper Account Team or Juniper Partner.

Features of 400ZR and 400G OpenZR+

IN THIS SECTION

- [Application Selection | 329](#)
- [Configure Application Selection | 332](#)
- [Target Output Power | 335](#)
- [Configure Target Output Power | 336](#)
- [Enhanced Loopback Options | 337](#)
- [Tunable Wavelength | 338](#)
- [Signal Power Performance | 338](#)
- [Ports in High-Power Mode | 341](#)
- [Configure 400ZR and 400G OpenZR+ Coherent Optics | 342](#)
- [Coherent Optics Performance Monitoring | 343](#)
- [User Configurable PM Interval Length | 345](#)
- [Configure User Configurable PM Interval Length | 346](#)

Application Selection

The 400ZR and 400G OpenZR+ optics advertises supported applications for a particular speed. You can select any of the application based on your requirement. You can also switch between these applications. See [Table 119 on page 329](#) for devices that support application selection.

Table 119: Supported Applications of ZR Optics

Application	Host ID	Host Interface Code Description	Media ID	Media Interface Code Description	Channelization
1	17	400GAUI-8 C2M (Annex 120E)	62	400ZR, DWDM, amplified	1x400G CFEC amplified

Table 119: Supported Applications of ZR Optics (Continued)

Application	Host ID	Host Interface Code Description	Media ID	Media Interface Code Description	Channelization
2	17	400GAUI-8 C2M (Annex 120E)	63	400ZR, single wavelength unamplified	1x400G CFEC unamplified
3	13	100GAUI-2 C2M (Annex 135G)	62	400ZR, DWDM, amplified	4x100G CFEC amplified

Table 120: Supported Applications of 400G OpenZR+ Optics

Application	Host ID	Host Interface Code Description	Media ID	Media Interface Code Description	Channelization
1	17	400GAUI-8 C2M (Annex 120E)	62	400ZR, DWDM, amplified	1x400G CFEC amplified
2	17	400GAUI-8 C2M (Annex 120E)	63	400ZR, Single Wavelength, unamplified	1x400G CFEC unamplified
3	13	100GAUI-2 C2M (Annex 135G)	62	400ZR, DWDM, amplified	4x100G CFEC amplified
4	17	400GAUI-8 C2M (Annex 120E)	70	ZR-400-OFEC-16QAM	1x400G OFEC
5	13	100GAUI-2 C2M (Annex 135G)	70	ZR-400-OFEC-16QAM	4x100G OFEC

Table 120: Supported Applications of 400G OpenZR+ Optics (Continued)

Application	Host ID	Host Interface Code Description	Media ID	Media Interface Code Description	Channelization
6	13	100GAUI-2 C2M (Annex 135G)	71	ZR-300-OFEC-8QAM	3x100G OFEC
7	13	100GAUI-2 C2M (Annex 135G)	72	ZR-200-OFEC-QPSK	2x100G OFEC
8	13	100GAUI-2 C2M (Annex 135G)	73	ZR-100-OFEC-QPSK	1x100G OFEC

Table 121: Default Application for 400ZR and 400G OpenZR+

Optics	Speed	Default Application
400GbE ZR	1x400G	CFEC amplified
400GbE ZR+		OFEC

For example:

To switch from 1x400G OFEC to 4x100G CFEC in ZR+

- Change the speed from 1x400 Gbps to 4x100 Gbps using the `set interface et-> number-of-sub-ports <: speed <:>`:

```
set interface et-> number-of-sub-ports 4 speed 100g
```

The command changes the speed from 1x400 Gbps to 4x100 Gbps.

- Switch the application from OFEC to CFEC using the `set interfaces <interface> optics-options application hostid <hostid> mediaid <mediaid> [domainid <domainid>]` command:

```
set interface et-<>:0 optics-options application mediaid 62 hostid 13
```

The command switches the application to 4x100 Gbps CFEC.

To display the configurable applications

- Use the `show interfaces diagnostics optics-applications <interface>` command to display the configurable applications. For example, 1x400G supports the applications as highlighted in the following output.

```
router> show interfaces diagnostics optics-applications et-0/0/10
Physical interface: et-0/0/10
Interface Name : et-0/0/10
Current Speed : 1x400G
Ap Sel Host Intf Code Media Intf Code Host Lanes Media Lanes Host Assign Media Assign
1 400GAUI-8 C2M (Annex 120E)(17) 400ZR, DWDM, amplified(62) 8 1 1 1
2 400GAUI-8 C2M (Annex 120E)(17) 400ZR, Single Wavelength,Unamplified(63) 8 1 1 1
3 100GAUI-2 C2M (Annex 135G)(13) 400ZR, DWDM, amplified(62) 2 1 85 1
```

Configure Application Selection

Execute the following steps to configure the 400ZR or 400G OpenZR+ optics:

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

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

2. Include the speed statement at the `[edit interfaces interface-name]` hierarchy level to configure the speed.

```
[edit interfaces interface-name]
user@host# set speed (400g | 100g);
```

3. Configure speed and channelization.

```
[edit interfaces interface-name]
user@host# set number-of-sub-ports <number-of-sub-ports>
```

To configure a different application for a speed, see [Step 5](#).

4. Display all configurable applications in the configured speed. You can use this information while executing [Step 5](#).

```
show interfaces diagnostics optics-applications <interface>
```

```
root@user> show interfaces diagnostics optics-applications et-0/0/10
Physical interface: et-0/0/10
Interface Name : et-0/0/10
Current Speed : 1x400G
Ap Sel Host Intf Code Media Intf Code Host Lanes Media Lanes Host Assign Media Assign
1 400GAUI-8 C2M (Annex 120E)(17) 400ZR, DWDM, amplified(62) 8 1 1 1
2 400GAUI-8 C2M (Annex 120E)(17) 400ZR, Single Wavelength,Unamplified(63) 8 1 1 1
3 100GAUI-2 C2M (Annex 135G)(13) 400ZR, DWDM, amplified(62) 2 1 85 1
```

5. Change the application.

a. Change to the [edit interfaces interface-name optics-options] hierarchy.

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

b. Execute the following command in this hierarchy:

```
set interfaces <interface> optics-options application hostid <hostid> mediaid <mediaid>
[ domainid <domainid>]
```

6. Display the output after you switch to a new application. The following output shows **media id** and **host id** as **62** and **17** respectively.

```
show interfaces et-0/0/10
```

```
Physical interface: et-0/0/10, Enabled, Physical link is Down
Interface index: 1046, SNMP ifIndex: 551
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None,
```

```
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Disabled, Media type: Fiber
Wavelength : 1550.12 nm, Frequency: 193.40 THz
Optic-loopback : Disabled Media Code : 400ZR, DWDM, amplified Host Code : 400GAUI-8 C2M
(Annex 120E)
Device flags : Present Running Down Transceiver-Rx-Alarm
Interface flags: Hardware-Down SNMP-Traps Internal: 0x80000
CoS queues : 8 supported, 8 maximum usable queues
Current address: 40:de:ad:8c:ca:22, Hardware address: 40:de:ad:8c:ca:22
Last flapped : 2022-12-12 02:53:51 PST (00:01:03 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Active alarms : LINK
Active defects : LINK, LOCAL-FAULT
PCS statistics Seconds
Bit errors 0
Errored blocks 0
Ethernet FEC Mode : FEC119
Ethernet FEC statistics Errors
FEC Corrected Errors 1275
FEC Uncorrected Errors 24
FEC Corrected Errors Rate 0
FEC Uncorrected Errors Rate 0
Optic FEC Mode : CFEC <-----
Optic FEC statistics:
Corrected Errors 65254593448
Uncorrected Words 0
Corrected Error rate 0
Uncorrected Error rate 0
Corrected Error Ratio ( 56 seconds average) 5.00e-01
PRBS Mode : Disabled
Interface transmit statistics: Disabled
Link Degrade :
Link Monitoring : Disable
Logical interface et-0/0/10.16386 (Index 1041) (SNMP ifIndex 606)
Flags: Up SNMP-Traps Encapsulation: ENET2 DF
Input packets : 0
Output packets: 0
Protocol multiservice, MTU: Unlimited
Flags: None
```

Target Output Power

You can view the minimum and maximum configurable transmit output power for 400ZR and 400G OpenZR+ optics modules. If you want to configure a transmit output power other than the default power, use the `set interfaces et-> optics-options tx-power <value>` command.

The minimum and maximum transmit output power varies:

- Based on the optical transceivers.
- Based on the transceiver vendor.

Table 122: Maximum and Minimum Tx Power

Optical Transceiver	JPN	Maximum Tx Power	Minimum Tx Power
QDD-400G-ZR	740-114884	4 dBm	-22 dBm
QDD-400G-ZR-M	740-131169	-10 dBm	-14 dBm
QDD-400G-ZR-M-HP	740-131168	1 dBm	-6 dBm
JCO400-QDD-ZR	740-157132	-10 dBm	-14 dBm
JCO400-QDD-ZR-M	740-157138	-10 dBm	-14 dBm
JCO400-QDD-ZR-M-HP	740-151745	1 dBm	-6 dBm

Changing or adjusting to a different Tx power improves transmission quality. See [tx-power](#).

Table 123: Default Output Power

Transceiver	Default Transmit Power
QDD-400G-ZR	-10 dBm
JCO400-QDD-ZR	-10 dBm
QDD-400G-ZR-M	-10 dBm

Table 123: Default Output Power (*Continued*)

Transceiver	Default Transmit Power
JCO400-QDD-ZR-M	-10 dBm
QDD-400G-ZR-M-HP	0 dBm

Configure Target Output Power

Execute the following steps to configure the 400ZR or 400G OpenZR+ optics:

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

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

2. Configure output power.

```
user@host# set interface et-<interface> optics-option tx-power <tx power>
```

3. Verify the target output power using the following command.

```
show interfaces diagnostics optics <interface>
```

```
root@user# run show interfaces diagnostics optics et-0/0/10
Physical interface: et-0/0/10
Module temperature : 41 degrees C / 105 degrees F
Module voltage : 3.18 V
Wavelength channel number : 28
Wavelength setpoint : 1550.12 nm
Tx dither : Disabled
Frequency error : 0.00 GHz
Wavelength error : 0.0 nm
TEC fault alarm : False
Wavelength unlocked alarm : False
Tx tune alarm : False
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 temperature high alarm threshold : 76 degrees C / 168 degrees F
Module temperature low alarm threshold : -6 degrees C / 21 degrees F
Module temperature high warning threshold : 73 degrees C / 163 degrees F
Module temperature low warning threshold : -3 degrees C / 26 degrees F
Module voltage high alarm threshold : 3.63 V
Module voltage low alarm threshold : 3.0 V
Module voltage high warning threshold : 3.599 V
Module voltage low warning threshold : 3.049 V
Laser bias current high alarm threshold : 524.28 mA
Laser bias current low alarm threshold : 0.0 mA
Laser bias current high warning threshold : 524.28 mA
Laser bias current low warning threshold : 0.0 mA
Laser output power high alarm threshold : 0.199 mW / -7.00 dBm
Laser output power low alarm threshold : 0.015 mW / -18.01 dBm
Laser output power high warning threshold : 0.158 mW / -7.99 dBm
Laser output power low warning threshold : 0.019 mW / -16.98 dBm
Laser rx power high alarm threshold : 2.511 mW / 4.00 dBm
Laser rx power low alarm threshold : 0.004 mW / -23.01 dBm
Laser rx power high warning threshold : 1.995 mW / 3.00 dBm
Laser rx power low warning threshold : 0.009 mW / -20.00 dBm
Laser temperature high alarm threshold : 76 degrees C / 168 degrees F
Laser temperature low alarm threshold : -6 degrees C / 21 degrees F
Laser temperature high warning threshold : 73 degrees C / 163 degrees F
Laser temperature low warning threshold : -3 degrees C / 26 degrees F
Lane 0
Laser bias current : 234.8 mA
Laser output power : 0.063 mW / -10.00 dBm ----->tx-power

```

Enhanced Loopback Options

Only the **Media Side Output** loopback type was available before Junos OS Evolved 23.1 R1.

From Junos OS Evolved 23.1 R1 onwards, the following types of loopbacks are available in addition to **Media Side Output**:

- Media Side Input—Loopback of the transceiver optical input signal back to its input
- Host Side Input—Loopback of the data packet signal output at the host side back to its input
- Host Side Output—Loopback of the data packet signal input at the host side back to its input

The enhanced command to configure `lo0` is `set interface et-<> optics-options loopback loopbacktype <loopbacktype>`.

For example:

```
set interface et-<> optics-options loopback loopbacktype <media-side-out>
set interface et-<> optics-options loopback loopbacktype <media-side-in>
set interface et-<> optics-options loopback loopbacktype <host-side-out>
set interface et-<> optics-options loopback loopbacktype <host-side-in>
```

This enhancement is supported on PTX10001-36MR. See [optics-options](#).

Tunable Wavelength

The 400ZR and 400G OpenZR+ are DWDM tunable transceivers. You can configure the transceivers with Optical Internetworking Forum (OIF) channels.

Both 400ZR and 400G OpenZR+ optic modules support flex grid channel spacing in increments of 6.25 GHz. The 400ZR and 400G OpenZR+ support 75 GHz/100 GHz frequency spacing. A 75 GHz channel spacing is more suitable for point-to-point links, while a 100 GHz channel spacing is more applicable to ROADM networks.

The `show optics diagnostics` command displays the tunable parameters like Thermoelectric Cooler (TEC) fault, wavelength unlocked fault status, wavelength/frequency, channel number, and so on.

Use the `set interface et-<> optics-options wavelength` command to configure wavelength. By default, the wavelength is 1550.12 nm, which corresponds to 193.40 THz. See [wavelength](#).

Signal Power Performance

From 23.4R1 Junos OS Evolved onwards, the `show interface diagnostics optics et-<x/x/x>` command shows the Rx signal power in the mW/dBm format. This enhancement is applicable to all devices on Junos OS Evolved platforms.

Example:

```
root@ardbeg-q> show interfaces diagnostics optics et-0/0/0
Physical interface: et-0/0/0
  Module temperature          : 55 degrees C / 131 degrees F
  Module voltage              : 3.252 V
  Wavelength channel number   : 28
  Wavelength setpoint         : 1550.12 nm
  Tx dither                  : Disabled
  Frequency error             : 0.00 GHz
  Wavelength error            : 0 nm
  TEC fault alarm             : False
  Wavelength unlocked alarm   : False
  Tx tune alarm               : False
  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 temperature high alarm threshold: 76 degrees C / 168 degrees F
  Module temperature low alarm threshold: -6 degrees C / 21 degrees F
  Module temperature high warning threshold: 73 degrees C / 163 degrees F
  Module temperature low warning threshold: -3 degrees C / 26 degrees F
  Module voltage high alarm threshold: 3.7 V
  Module voltage low alarm threshold: 3 V
  Module voltage high warning threshold: 3.599 V
  Module voltage low warning threshold: 3.049 V
  Laser bias current high alarm threshold: 120 mA
  Laser bias current low alarm threshold: 80 mA
  Laser bias current high warning threshold: 110 mA
  Laser bias current low warning threshold: 90 mA
  Laser output power high alarm threshold: 2.511 mW / 4.00 dBm
  Laser output power low alarm threshold: 0.125 mW / -8.99 dBm
  Laser output power high warning threshold: 1.995 mW / 3.00 dBm
  Laser output power low warning threshold: 0.158 mW / -7.99 dBm
  Laser rx power high alarm threshold: 2.511 mW / 4.00 dBm
  Laser rx power low alarm threshold: 0.006 mW / -22.00 dBm
  Laser rx power high warning threshold: 1.995 mW / 3.00 dBm
  Laser rx power low warning threshold: 0.012 mW / -18.99 dBm
```

Laser temperature high alarm threshold	:	65 degrees C / 149 degrees F
Laser temperature low alarm threshold	:	25 degrees C / 77 degrees F
Laser temperature high warning threshold	:	55 degrees C / 131 degrees F
Laser temperature low warning threshold	:	35 degrees C / 95 degrees F
Lane 0		
Laser bias current	:	100 mA
Laser output power	:	1.004 mW / 0.01 dBm
Laser temperature	:	48 degrees C / 118 degrees F
Laser receiver power	:	0.887 mW / -0.51 dBm
Rx power (signal)	:	0.00 mW / -44.00 dBm <-----
Lane chromatic dispersion	:	0.0 ps/nm
Lane differential group delay	:	0.0 ps
Lane carrier frequency offset	:	0.0 MHz
Lane polarization dependent loss	:	0.0 dB
Lane snr	:	0.0 dB
Lane Optical signal-to-noise ratio	:	0.0 dB
Lane sopmd	:	0.0
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 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
Laser output power high alarm	:	Off
Laser output power low alarm	:	Off
Laser output power high warning	:	Off
Laser output power low warning	:	Off
Tx loss of signal functionality alarm	:	Off
Rx loss of signal alarm	:	Off
Tx laser disabled alarm	:	Off

Ports in High-Power Mode

The number of ports supporting the 400G-ZR optics is restricted based on the power consumption on the QFX5220-32CD and QFX5130-32CD devices. For better thermal handling, 16 ports out of 32 QSFP56-DD ports (0, 3, 4, 7, 8, 11, 12, 15, 16, 19, 20, 23, 24, 27, 28, 31) support the 400G-ZR optics.

The below table includes the corresponding ports that are configured as “unused”.

Table 124: Ports Supporting 400G-ZR Optics

Ports Supporting 400G-ZR Optics	Corresponding Ports to Be Configured to Unused
0	1
3	2
4	5
7	6
8	9
11	10
12	13
15	14
16	17
19	18
20	21
23	22
24	25
27	26
28	29

Table 124: Ports Supporting 400G-ZR Optics (Continued)

Ports Supporting 400G-ZR Optics	Corresponding Ports to Be Configured to Unused
31	30

See No Link Title to configure the ports in the high-power mode.

Configure 400ZR and 400G OpenZR+ Coherent Optics

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

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

2. Include the speed statement at the [edit interfaces interface-name] hierarchy level to configure the speed.

```
[edit interfaces interface-name]
user@host# set speed (400g | 100g);
```

3. Configure speed and channelization.

```
[edit interfaces interface-name]
user@host# set number-of-sub-ports <number-of-sub-ports>
```

You must use [edit chassis fpc *fpc-slot* pic *pic-number*] set number-of-sub-ports *number-of-sub-ports* hierarchy to configure speed and channelization for PTX10003 and MX Series routers.

4. Configure a different application for a speed.
 - a. Change to the [edit interfaces interface-name optics-options].

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

b. Execute the following command in this hierarchy.

```
user@host# set interfaces <interface> optics-options application hostid <hostid> mediaid <mediaid> [domainid <domainid>]
```

5. [Optional] Configure output power in the [edit interfaces interface-name optics-options] hierarchy level. If in channelized mode, you must use the first sub-port to configure output power.

```
user@host# set interface et-<interface> optics-option tx-power <tx power>
```

Example:

```
user@host# set interfaces et-0/0/0:0 optics-options tx-power -10
```

6. Specify the wavelength of optics in nm. If in channelized mode, you must use the first sub-port to configure wavelength.

```
[edit interfaces interface-name optics-options]  
user@host# set wavelength nm
```

Example:

```
user@host# set interfaces et-0/0/0:0 optics-options wavelength 1552.52
```

Coherent Optics Performance Monitoring

SUMMARY

Learn about performance monitoring (PM) for 400ZR and 400G OpenZR+ coherent optics.

IN THIS SECTION

- Benefits of Coherent Optics Performance Monitoring | [344](#)
- WHAT's NEXT | [345](#)

Coherent optics performance monitoring (PM) stores diagnostic data for optical parameters over defined intervals. PM offers insights into optical components' performance and health, allowing proactive maintenance and troubleshooting. PM helps users understand data collection intervals, statistics types (minimum, maximum, average), and data access. Users can monitor short-term and long-term performance trends, supporting efficient network management.

- **Interval Management:** The PM manages data collection in 15-minute and 1-day intervals. The 15-minute intervals begin at the top of the hour, 15, 30, and 45 minutes past the hour, while the 1-day intervals begin at midnight UTC.
- **Data Retention:** The system retains data for up to 96 historical 15-minute intervals and one historical 1-day interval, providing a rolling 24-hour coverage and the previous full day's data.
- **Continuous Data Update:** The PM continuously updates the minimum, maximum, and average values for each monitored parameter based on data samples polled every 1-5 seconds.

Frequent polling and storing allow you to monitor real-time performance and historical trends of optical components. This provides crucial insights into your optical network's overall health and performance.

The PM feature supports Threshold Crossing Alerts (TCA). TCA immediately inform the operator when any monitored performance parameter crosses a defined threshold. TCA warns the operator about possible deterioration in the optical module or its traffic.

To display PM and TCA data

Use the `show interfaces transport pm` command to display the current and historical performance metrics. See [show interfaces transport pm](#).

Benefits of Coherent Optics Performance Monitoring

- **Proactive Maintenance:** PM continuously collects and analyzes performance data. This process helps identify potential issues early, enabling timely maintenance and reducing the risk of network outages.
- **Accurate Performance Metrics:** Monitor minimum, maximum, and average values to ensure precision and get accurate view of optical component performance over time.
- **Efficient Troubleshooting:** Access real-time and historical performance data to quickly identify and resolve network issues.
- **Enhanced Network Management:** PM offers vital insights into optical components' health and performance. This action boosts decision-making and resource optimization, resulting in better network management and operation.

WHAT's NEXT

Learn about user configurable performance monitoring (PM) interval lengths for more granular and precise network performance monitoring. See ["User Configurable PM Interval Length" on page 345](#).

User Configurable PM Interval Length

SUMMARY

Learn about user configurable performance monitoring (PM) interval lengths for more granular and precise network performance monitoring.

IN THIS SECTION

- [Benefits of User Configurable PM Interval Length | 345](#)
- [WHAT's NEXT | 346](#)

User configurable performance monitoring (PM) interval lengths provide enhanced granularity in performance monitoring within your networking system. Adjust the default 15-minute PM intervals to shorter durations such as 10 seconds, 30 seconds, 1 minute, or 5 minutes for more precise network performance monitoring.

Synchronizing PM intervals with system time ensures predictable and consistent data collection. For instance, the system triggers 10-second intervals at the start of every minute and repeat at 10, 20, 30, 40, and 50 seconds past each minute. This synchronization helps anticipate PM bin resets and data accumulation, simplifying performance trend analysis.

Modifying the PM interval length deletes existing PM bins from the previous configuration. Be aware of the implications for ongoing data collection and historical data integrity. Synchronizing intervals with system time and deleting existing bins preserve data accuracy and consistency, enabling effective network performance monitoring.

Benefits of User Configurable PM Interval Length

- **Precise Network Performance:** Provides enhanced granularity in performance monitoring, allowing you to capture more detailed and precise network performance data.
- **Shorter Monitoring Intervals:** Facilitates quicker identification and resolution of network issues with shorter monitoring intervals, leading to improved network reliability and uptime.
- **Synchronized Intervals:** Ensures predictable data collection and analysis with synchronized intervals, helping you consistently track performance metrics.

- **Historical Performance data:** Maintains historical performance data integrity by supporting up to 96 historical PM bins, even when the interval lengths are modified.

WHAT's NEXT

Learn how to configure user configurable performance monitoring (PM) interval lengths for more granular and precise network performance monitoring. See "[Configure User Configurable PM Interval Length](#)" on page 346.

SEE ALSO

[interval-length \(Chassis\)](#)
[optics \(Chassis\)](#)

Configure User Configurable PM Interval Length

SUMMARY

Learn how to configure user configurable performance monitoring (PM) interval lengths for more granular and precise network performance monitoring.

Use [400ZR and 400G OpenZR+](#) to confirm platform and release support.

Execute the following steps to configure the user configurable PM interval length for coherent optics:

1. In configuration mode, go to the [edit chassis optics pm] hierarchy level.

```
[edit]
user@host# edit chassis optics pm
```

2. Include the interval length at the [edit chassis optics pm] hierarchy level to configure the user configurable PM interval length.

```
[edit chassis optics pm]
user@host# set interval-length (10s | 30s | 1min | 5min | 15min)
```

Example:

```
user@host# set chassis optics pm interval-length 5min
```

3. To view the current and historical performance metrics, use the `show interfaces transport pm` command. See [show interfaces transport pm](#).

The command applies globally to all ports with coherent optics supporting PM. Users must configure the PM interval length based on monitoring needs or network conditions.

When the configuration is deleted, the system reverts the interval length back to the default 15-minute duration.

SEE ALSO

[interval-length \(Chassis\)](#)
[optics \(Chassis\)](#)

Configure and Monitor 400ZR and 400G OpenZR+ Coherent Optics

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

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

2. Include the `speed` statement at the `[edit interfaces interface-name]` hierarchy level to configure the speed.

```
[edit interfaces interface-name]  
user@host# set speed (400g | 100g);
```

3. Configure speed and channelization.

```
[edit interfaces interface-name]
user@host# set number-of-sub-ports <number-of-sub-ports>
```

You must use [edit chassis fpc *fpc-slot* pic *pic-number*] set number-of-sub-ports *number-of-sub-ports* hierarchy to configure speed and channelization.

4. Configure a different application for a speed.

a. Change to the [edit interfaces interface-name optics-options].

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

b. Execute the following command in this hierarchy.

```
user@host# set interfaces <interface> optics-options application hostid <hostid> mediaid
<mediaid> [domainid
<domainid>]
```

5. [Optional] Configure output power in the [edit interfaces interface-name optics-options] hierarchy level.
If in channelized mode, you must use the first sub-port to configure output power.

```
user@host# set interface et-<interface> optics-option tx-power <tx power>
```

Example:

```
user@host# set interfaces et-0/0/0:0 optics-options tx-power -10
```

6. Specify the wavelength of optics in nm. If in channelized mode, you must use the first sub-port to configure wavelength.

```
[edit interfaces interface-name optics-options]
user@host# set wavelength nm
```

Example:

```
user@host# set interfaces et-0/0/0:0 optics-options wavelength 1552.52
```

7. To view the current and historical performance metrics, use the `show interfaces transport pm` command. See [show interfaces transport pm](#).

7

PART

100GbE, 40GbE, and 10GbE Optical Transceivers

- 100GbE, 40GbE, and 10GbE | 351

100GbE, 40GbE, and 10GbE

IN THIS SECTION

- [Overview | 351](#)
- [Supported Features | 353](#)
- [Supported PICs | 357](#)
- [Configure Optics | 358](#)
- [Supported OTN and Optics Options | 363](#)
- [Forward Error Correction and Bit Error Rate | 400](#)
 - [| 412](#)
 - [| 412](#)
 - [| 412](#)
 - [| 412](#)
 - [| 412](#)
- [Interface Mapping and Modulation format for PTX10K-LC1104 Line Card | 413](#)

Read this topic for information about optics support that the specific line cards and devices provide.

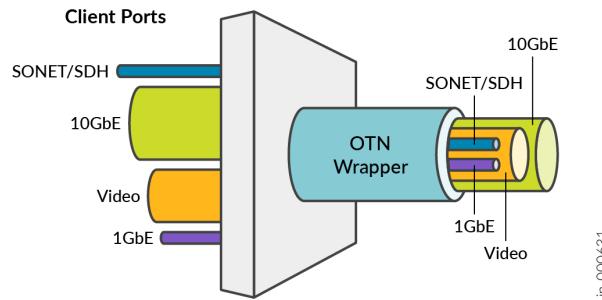
Overview

IN THIS SECTION

- [Benefits | 352](#)

Optical transmission leverages properties of light waves, including amplitude, phase, and polarization to optimize the capacity of a fiber optic link.

Optics supports Optical Transport Network (OTN), a standard defined by ITU G.709. The standard defines OTN as a set of optical network elements that are connected by optical fiber links.



Optics performs the following functions on optical channels that carry client signals:

- Transport
- Multiplexing
- Switching
- Management
- Supervision

Optics bridges the gap between next-generation IP and legacy TDM networks. When you add optics to a Dense Wavelength Division Multiplexing (DWDM) network, you can easily manage a channel. For example, you can configure, troubleshoot, or see alarms. It is a standardized method for transparent transport of services over optical wavelengths. Hence, modification of the client signal does not occur when it passes through the DWDM network.

ITU G.709 standard ensures:

- Performance monitoring and alarm management as specified in ITU standard.
- Transparent transport of Ethernet signals with optical channel data unit 2 (ODU2) and ODU2e framing on a per-port basis.
- Bit error rate (BER) based on pre-forward error correction (pre-FEC). Fast reroute (FRR) uses the pre-FEC BER as an indication of the an OTN link condition.

Benefits

The OTN provides the following benefits:

- Enables FEC and thus improves system performance.
- Enables alarm handling and thus improves system monitoring.

- Protects the network from mixing of heterogeneous services, saves bandwidth, and transfers any traffic transparently.
- Supports dedicated Ethernet links at 1GbE, 10GbE, 40GbE, and 100GbE and thus improves scalability.

Supported Features

Table 125: Supported optics features

Features	Description
Interface	<p>Supports:</p> <ul style="list-style-type: none">• Minimum channel spacing of 6.25GHz and compliant with ITU G.709• FEC statistics, generic FEC, ultra-forward FEC, enhanced FEC, and no-FEC modes of operation• GbE local loopback, External clock.• Diagnostic tools: Line loopback and local loopback.

Table 125: Supported optics features (*Continued*)

Features	Description
	<p>Attributes:</p> <ul style="list-style-type: none"> • The interface names start with prefixes et and ot. • OTN payload pseudorandom binary sequence (PRBS) generation and checking by enabling or disabling PRBS. • Link-level pause frames can halt the Ethernet interface from transmitting packets for a configured period. • SNMP management of the interface based on RFC 3591, Definitions of Managed Objects for the Optical Interface Type: <ul style="list-style-type: none"> • Set functionality • Juniper Networks Black-Link MIB • Optics MIB • FRU MIB <p>Performance monitoring (15-minute and 1-day performance monitoring and historic statistics):</p> <ul style="list-style-type: none"> • Near-end and far-end • TCA • BER • FEC • Optical

Table 125: Supported optics features (*Continued*)

Features	Description
L2 and L3 Features	<p>Supports:</p> <ul style="list-style-type: none"> • MAC statistics, MPLS fast reroute, SNMP, and Flow control. • MAC oversized packet counters based on default MTU value or user-configured MTU value. • Per-port source address MAC filter, per-destination address MAC filter, per-physical interface source address MAC filter, per-logical interface source address MAC accounting, and source MAC filter per physical interface. • 802.3 ah OAM and 802.1 ag OAM. • Logical interfaces on an aggregated Ethernet physical interface. • VLAN tagging, flexible VLAN tagging, per-VLAN queuing (using Packet Forwarding Engine), and S-VLAN tagging.
Threshold Crossing Alarms	<p>You receive TCA when you cross a certain configurable threshold (near-end measurement threshold or far-end measurement threshold). The TCA remains until the end of the 15-minute interval for parameters such as OTU and ODU</p>
	<p>Supported alarms:</p> <ul style="list-style-type: none"> • Background block error threshold (BBE) • Errored seconds (ES) threshold • Severely errored second (SES) threshold • Unavailable seconds (UAS) threshold

OTN Alarms and Defects

Table 126: OTN Alarms and Defects

OTN Alarms and Defects	Description
CSF	Client Signal Failure
LOS	Loss of signal
LOF	Loss of frame
LOM	Loss of multi frame
SSF	Server Signal Failure
TSF	Trail Signal Fail
OTU-FEC-DEG	FEC Degraded
OTU-FEC-EXE	Excessive Errors, FEC_FAIL from the transponder
OTU-AIS	Alarm Indication Signal
OTU-BDI	Backward Defect Identification
OTU-IAE	Incoming Alignment Error
OTU-TTIM	Destination Access Point Identifier [DAPI], Source Access Point Identifier [SAPI], or both mismatch from expected to be received.
OTU-SD	Signal Degrade
OTU-SF	Signal Fail
ODU-LCK	ODU lock triggers for PM [path monitoring]

Table 126: OTN Alarms and Defects (*Continued*)

OTN Alarms and Defects	Description
ODU-AIS	Alarm indication signal
ODU-OCI	Open connection indication
ODU-BDI	Backward defect indication
ODU-IAE	Incoming alignment error
ODU-DAPI-TTIM	DAPI or DAPI/SAPI mismatch from expected to receive.
ODU-SAPI-TTIM	SAPI or DAPI/SAPI mismatch from expected to receive.
ODU-BEI	Backward Error Indication
ODU-SSF	Server Signal Fail
ODU-TSF	Trail Signal Fail
ODU-SD	Signal Degrade.
ODU-SF	Signal Fail
OPU-PTM	Payload Type Mismatch.

Supported PICs

Table 1 describes PICs that support optics.

Table 127: Supported PICs

PIC	
P1-PTX-24-10G-W-SFPP	PTX3000: Junos OS Release 13.2R2 and later PTX5000: Junos OS Release 12.3R2 and later Junos OS Release 13.2R1 and later
P2-100GE-OTN	Junos OS Release 15.1F6 Junos OS Release 16.1R2 and later Junos OS Release 17.1R1 and later
MIC3-100G-DWDM	Junos OS Release 15.1F5 and 15.1F6 Junos OS Release 17.1R1 and later
PTX-5-100G-WDM	PTX3000: Junos OS Release 15.1F6, Junos OS Release 17.1R1 and later PTX5000: Junos OS Release 15.1F6, Junos OS Release 17.1R1 and later

Configure Optics

This topic provides information about how to configure optics interface, OTN options on an interface, and optics options on an interface.

Configure Interfaces.

To configure the interface-specific options:

1. Go to the [edit interface *interface-name*] hierarchy level, where *interface-name* is in the et-fpc/pic/port format.

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

2. Configure the VLAN tagging option on the OTN interface to enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

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

3. Configure the maximum transmission unit (MTU) size in bytes for the interface.

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

4. Configure a VLAN ID for the interface.

```
[edit interfaces interface-name]
user@host# set vlan-id number
```

5. Configure the family for the interface.

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

6. Configure an IP address for the interface.

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

Configure OTN Options on the interface

To configure the OTN-related options on the interface:

1. Go to the [edit interface *interface-name* otn-options] hierarchy level:

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

2. Enable the OTN mode as OTU2e, OTU1e, or OTU2 for the interface.

```
[edit interfaces interface-name otn-options]
user@host# set rate fixed-stuff-bytes|no-fixed-stuff-bytes|oc192
```



NOTE: fixed-stuff-bytes is for OTU2e rate, no-fixed-stuff-bytes is for OTU1e rate and oc192 is for OTU2 rate. OTU2e and OTU1e rates are applicable for LAN PHY framing mode. OTU2 is applicable for WAN PHY framing mode. Framing mode is to set through the set interfaces framing configuration statement.

3. Enable the laser on the OTN interface. The laser is disabled by default for all OTN interfaces.

```
[edit interfaces interface-name otn-options]
user@host# set laser-enable
```

4. Set a trail trace identifier for the source access point and for the destination access point for ODU and OTU on the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set tti (odu-dapi | odu-expected-receive-dapi | odu-expected-receive-sapi | odu-sapi | otu-dapi | otu-expected-receive-dapi | otu-expected-receive-sapi | otu-sapi) tti-identifier
```

5. Ignore the trigger for the defect or set the hold time.

Configure the hold time for the defect trigger as:

- *up* with a value—Wait for the hold time delay before clearing the alarm when the defect is absent on the OTN interface.
- *down* with a value—Wait for the hold time delay before raising the alarm when the defect occurs for the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set trigger (oc-lof | oc-lom | oc-los | oc-tsf | odu-ais | odu-bdi | odu-bei | oduiae | odu-lck | odu-oci | odu-sd | odu-ttim | opu-ptim | otu-ais | otu-bdi | otu-fec-deg | otu-fec-exe | otuiae | otu-sd | otu-ttim) (hold-time (down value | up value) | ignore)
```

6. Enable the threshold crossing alarms for the OTN interface along with the trigger for the defect.

```
[edit interfaces interface-name otn-options]
user@host# set tca (odu-tca-bbe | odu-tca-es | odu-tca-ses | odu-tca-uas | otu-tca-bbe | otu-tca-es | otu-tca-ses | otu-tca-uas ) (enable-tca | no-enable-tca | threshold)
```

7. Set the OTN header bytes as a transmit payload type from 0 bytes through 255 bytes for the packets that are transmitted on the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set bytes transmit-payload-type value
```

8. Configure the forward error correction (FEC) mode as Generic Forward Error Correction (GFEC), Enhanced Forward Error Correction (EFEC), Ultra Forward Error Correction (UFEC), or no-FEC (none) for the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set fec (gfec | ufec | efec | none)
```

9. Enable a consequent action as listed in the ITU-T G.798 standard for ODU trail trace identifier mismatch (TTIM) on the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set odu-ttim-action-enable
```

10. Enable a consequent action as listed in the ITU-T G.798 standard for OTU trail trace identifier mismatch (TTIM) on the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set otu-ttim-action-enable
```

11. Configure the threshold value for signal degradation when an alarm needs to be raised. Configure the threshold value after signal degradation when the alarm needs to be cleared. When you configure the interval along with the ber-threshold-signal-degrade *value* statement, the bit error rate (BER) must remain above the signal degradation threshold for the configured interval after which the alarm is raised. When the interval is configured along with the ber-threshold-clear *value*

statement, then BER must remain below the clear threshold for the configured interval after which the alarm is cleared.

```
[edit interfaces interface-name otn-options signal-degrade]
user@host# set ber-threshold-signal-degrade value
user@host# set ber-threshold-clear value
user@host# set interval value
```

12. Enable the following actions for the preemptive-fast-reroute statement:

- Backward FRR—Insert the local pre-FEC status into the transmitted OTN frames and monitor the received OTN frames for the pre-FEC status.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set backward-frr-enable
```

- Monitoring of signal degradation of pre-FEC OTN frames.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set signal-degrade-monitor-enable
```

Configure Optics Options on the interface

To configure the optics-specific options on the interface:

1. Specify the modulation format at the [edit interface *interface-name* optics-options] hierarchy level.

```
[edit interfaces interface-name optics-options]
user@host# set modulation-format (qpsk/8qam/16qam)
```

2. Specify encoding.

```
[edit interfaces interface-name optics-options]
user@host# set encoding (differential/non-differential)
```

3. Specify the optical transmit laser output power in dBm at the [edit interface *interface-name* optics-options] hierarchy level. The default transmit laser output value is 0 dBm.

```
[edit interfaces interface-name optics-options]
user@host# set tx-power value
```

4. Specify the wavelength of the optics in nanometers.

```
[edit interfaces interface-name optics-options]
user@host# set wavelength nm
```

SEE ALSO

[Supported PICs | 357](#)

optics-options

otn-options

signal-degrade

preemptive-fast-reroute

wavelength

Supported OTN and Optics Options

IN THIS SECTION

- [Supported OTN Options on ACX6360 and ACX5448 Routers | 364](#)
- [Supported OTN Options on MX Series Routers | 369](#)
- [Supported OTN Options on PTX Series Routers | 379](#)
- [Supported Optics Options on ACX6360 and ACX5448-D Routers | 391](#)
- [Supported Optics Options on PTX10008 and PTX10016 Series Routers | 395](#)

Read this topic for information about the supported optics options and OTN options on specific devices.

Supported OTN Options on ACX6360 and ACX5448 Routers

Table 128 on page 364 lists the statements that are supported on Juniper Networks® ACX6360 and ACX5448 routers at the [edit interfaces *interface-name* otn-options] hierarchy level. Note that the term *NA* denotes that the statement is not applicable for that particular component:

Table 128: Statements Supported on ACX6360 and ACX5448 Routers

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
bytes (otn-options)	transmit-payload- type <i>value</i>	Yes	Yes
insert- odu-lck	-	Yes	Yes
insert- odu-oci	-	Yes	Yes
is-ma no-is-ma	-	Yes	Yes
line-loopback no-line-loopback	-	Yes	Yes
local-loopback no-local-loopback	-	Yes	Yes
odu-ttim-action- enable no-odu-ttim-action- enable	-	Yes	Yes
otu-ttim-action- enable no-otu-ttim-action- enable	-	Yes	Yes
prbs no-prbs	-	Yes	Yes
preemptive-fast-reroute	backward-frr -enable no-backward-frr- enable	Yes	Yes

Table 128: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
tca	signal-degrade-monitor-enable no-signal-degrade-monitor-enable	Yes	Yes
	odu-backward-frr-enable no-odu-backward-frr-enable	No	No
	odu-signal-degrade-monitor-enable no-odu-signal-degrade-monitor-enable	No	No
	odu-tca-bbe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-bbe-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-es (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-es-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-ses (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes

Table 128: Statements Supported on ACX6360 and ACX5448 Routers (*Continued*)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
	odu-tca-ses-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-uas (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-bbe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-bbe-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-es (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-es-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-ses (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-ses-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes

Table 128: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
	otu-tca-uas (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-uas-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes
trigger <i>trigger-identifier</i>	oc-lof (hold-time (down up) ignore)	Yes	Yes
	oc-lom (hold-time (down up) ignore)	Yes	Yes
	oc-tsf (hold-time (down up) ignore)	Yes	Yes
	odu-ais (hold-time (down up) ignore)	Yes	Yes
	odu-bdi (hold-time (down up) ignore)	Yes	Yes
	odu-bei (hold-time (down up) ignore)	Yes	Yes
	odu-iae (hold-time (down up) ignore)	Yes	Yes
	odu-lck (hold-time (down up) ignore)	Yes	Yes

Table 128: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
	odu-oci (hold-time (down up) ignore)	Yes	Yes
	odu-sd (hold-time (down up) ignore)	Yes	Yes
	odu-ttim	Yes	Yes
	opu-ptim (hold-time (down up) ignore)	Yes	Yes
	otu-ais (hold-time (down up) ignore)	Yes	Yes
	otu-bdi (hold-time (down up) ignore)	Yes	Yes
	otuiae (hold-time (down up) ignore)	Yes	Yes
	otu-sd (hold-time (down up) ignore)	Yes	Yes
	otu-ttim (hold-time (down up) ignore)	Yes	Yes
tti <i>tti-identifier</i>	odu-dapi <i>identifier</i>	Yes	Yes
	odu-expected- receive-dapi <i>identifier</i>	Yes	Yes

Table 128: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
	odu-expected- receive-sapi <i>identifier</i>	Yes	Yes
	odu-sapi <i>identifier</i>	Yes	Yes
	otu-dapi <i>identifier</i>	Yes	Yes
	otu-expected- receive-dapi <i>identifier</i>	Yes	Yes
	otu-expected- receive-sapi <i>identifier</i>	Yes	Yes
	otu-sapi <i>identifier</i>	Yes	Yes

SEE ALSO

Configuring OTN

Supported OTN Options on MX Series Routers

[Table 129](#) on page 370 lists the statements that are supported on 100-Gigabit Ethernet MICs on Juniper Networks® MX Series Universal Routers at the [edit interfaces *interface-Name* otn-options] hierarchy level.

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
bytes	transmit-payload-type <i>value</i>	No	Yes
fec	(efec / gfec / gfec-sdfec / hgfec/ none / dsfec/ ufec)	Yes (ufec)	Yes (gfec,hgfec,sdfec)
insert- odu-lck	-	Yes	Yes
insert- odu-oci	-	Yes	Yes
is-ma no-is-ma	-	Yes	Yes
laser-eNoble no- laser-eNoble	-	Yes	Yes
line-loopback no-line-loopback	-	Yes	Yes
local-loopback no-local-loopback	-	Yes	Yes
odu-delay- maNogement	bypass no-bypass	No	Yes
	monitor- end-point no- monitor-end-point	No	Yes
	number- of-frames <i>value</i>	No	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	no-start- measurement start- measurement	No	Yes
sigNol-degrade	ber-threshold -clear <i>value</i>	Yes	Yes
	ber-threshold -sigNol- degrade <i>value</i>	Yes	Yes
	interval <i>value</i>	Yes	Yes
odu-ttim-action- eNoble no-odu- ttim-action- eNoble	-	Yes	Yes
otu-ttim-action- eNoble no-otu- ttim-action- eNoble	-	Yes	Yes
prbs no-prbs	-	Yes	Yes
preemptive-fast- reroute	backward-frr -eNoble no- backward-frr- eNoble	Yes	Yes
	sigNol- degrade- monitor- eNoble no-sigNol- degrade- monitor-eNoble	Yes	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	odu-backward-frr-eNoble no-odu-backward-frr-eNoble	No	Yes
	odu-sigNol-degrade- monitor-eNoble no-odu- sigNol-degrade- monitor- eNoble	No	Yes
rate	fixed-stuff-bytes no- fixed-stuff-bytes	Yes	Yes
	oc192	Yes	Yes
	otu4	Yes	Yes
	pass-through	Yes (pass-through)	Yes (pass-through)
sigNol-degrade	ber-threshold-clear <i>value</i>	Yes	Yes
	ber-threshold-sigNol- degrade <i>value</i>	Yes	Yes
	interval <i>value</i>	Yes	Yes
	q-threshold-sigNol-degrade	No	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	q-threshold-sigNol-degrade-clear	No	Yes
tca	odu-tca-bbe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-bbe-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-es (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-es-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-ses (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-ses-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	odu-tca-uas (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	otu-tca-bbe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-bbe-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-es (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-es-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-fec-ber (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-ses (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-ses-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
	otu-tca-uas (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	otu-tca-uas-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
<i>transport-monitoring</i>	-	No	Yes
trigger <i>trigger-identifier</i>	oc-lof (hold-time (down up) ignore)	Yes	Yes
	oc-lom (hold-time (down up) ignore)	Yes	Yes
	oc-los (hold-time (down up) ignore)	Yes	Yes
	oc-tsf (hold-time (down up) ignore)	No	Yes
	oc-wavelength-lock (hold- time (down up) ignore)	No	Yes
	odu-ais (hold-time (down up) ignore)	Yes	Yes
	odu-bdi (hold-time (down up) ignore)	Yes	Yes
	odu-bei (hold-time (down up) ignore)	Yes	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	odu-iae (hold-time (down up) ignore)	Yes	Yes
	odu-lck (hold-time (down up) ignore)	Yes	Yes
	odu-oci (hold-time (down up) ignore)	Yes	Yes
	odu-sd (hold-time (down up) ignore)	Yes	Yes
	odu-tca-es	Yes	No
	odu-tca-ses	Yes	No
	odu-tca-uas	Yes	No
	odu-ttim	Yes	Yes
	opu-ptim (hold-time (down up) ignore)	Yes	Yes
	otu-ais (hold-time (down up) ignore)	Yes	Yes
	otu-bdi (hold-time (down up) ignore)	Yes	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	otu-fec-deg (hold-time (down up) ignore)	No	Yes
	otu-fec-exe (hold-time (down up) ignore)	No	Yes
	otu-iae (hold-time (down up) ignore)	Yes	Yes
	otu-sd (hold-time (down up) ignore)	Yes	Yes
	odu-tca-es	Yes	No
	odu-tca-ses	Yes	No
	odu-tca-uas	Yes	No
	otu-ttim (hold-time (down up) ignore)	Yes	Yes
tti <i>tti-identifier</i>	odu-dapi <i>identifiery</i>	Yes	Yes
	odu-dapi-first-byte-nul no-odu-dapi-first-byte-nul	No	Yes
	odu-expected-receive- dapi <i>identifier</i>	Yes	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	odu-expected-receive-dapi-first-byte-nul no-odu-expected-receive-dapi-first-byte-nul	No	Yes
	odu-expected-receive-sapi <i>identifier</i>	Yes	Yes
	odu-sapi <i>identifier</i>	Yes	Yes
	odu-sapi-first-byte-nul no-odu-sapi-first-byte-nul	No	Yes
	otu-dapi <i>identifier</i>	Yes	Yes
	otu-dapi-first-byte-nul no-otu-dapi-first-byte-nul	No	Yes
	otu-expected- receive-dapi <i>identifier</i>	Yes	Yes
	otu-expected-receive-dapi-first-byte-nul no-odu-expected-receive-dapi-first-byte-nul	No	Yes
	otu-expected- receive-sapi <i>identifier</i>	Yes	Yes

Table 129: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

Statement	Options	MIC6-100G-CFP2 (MX2010 / MX2020) (13.3R3)	MIC3-100G-DWDM (MX240, MX480, MX960, MX2010, and MX2020) (15.1F5)
	otu-expected-receive-sapi-first-byte-nul no-odu-expected-receive-dapi-first-byte-nul	No	Yes
	otu-sapi <i>identifier</i>	Yes	Yes
	otu-sapi-first-byte-nul	No	Yes

SEE ALSO

Configuring OTN

Supported OTN Options on PTX Series Routers

[Table 130 on page 379](#) lists the statements that are supported on 100-Gigabit Ethernet PICs on Juniper Networks® PTX Series Routers at the [edit interfaces *interface-name* otn-options] hierarchy level.

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
bytes (otn-options)	transmit-payload-type <i>value</i>	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE- OTN (PTX5000) (14.1R2 / 14.2R1)	P1- PTX-24-1 0G-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
fec	(efec / gfec / gfec-sdfec / none / ufec)	Yes	Yes (gfec, none)	Yes	Yes
insert-odu- lck	-	Yes	Yes	Yes	Yes
insert-odu- oci	-	Yes	Yes	Yes	Yes
is-ma no- is-ma	-	Yes	NA	Yes	Yes
laser-enable no-laser- enable	-	Yes	Yes	Yes	Yes
line-loopback no-line- loopback	-	Yes	Yes	Yes	Yes
local- loopback no-local- loopback	-	Yes	Yes	Yes	Yes
odu-delay- management	bypass no-bypass	Yes	No	No	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	monitor- end-point no-monitor-end-point	Yes	No	No	Yes
	number- of-frames <i>value</i>	Yes	No	No	Yes
	no-start- measurement start- measurement	Yes	No	No	Yes
odu-signal-degrade	ber-threshold - clear <i>value</i>	No No	Yes	No	Yes
	ber-threshold - signal-degrade <i>value</i>	No No	Yes	No	Yes
	interval <i>value</i>	No	Yes	No	Yes
odu-ttim-action-enable no-odu-ttim-action-enable	-	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
otu-ttim-action-enable no-otu-ttim-action-enable	-	Yes	Yes	Yes	Yes
prbs no-prbs	-	Yes	Yes	Yes	Yes
preemptive-fast-reroute	backward-frr -enable no-backward-frr-enable	Yes	Yes	Yes	Yes
	signal-degrade-monitor-enable no-signal-degrade-monitor-enable	Yes	Yes	Yes	Yes
	odu-backward-frr-enable no-odu-backward-frr-enable	No	Yes	No	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	odu-signal-degrade-monitor-enable no-odu-signal-degrade-monitor-enable	No	Yes	No	Yes
rate	fixed-stuff-bytes no-fixed-stuff-bytes	Yes	No	Yes	Yes
	oc192	Yes	No	Yes	Yes
	otu4	Yes	Yes	No	Yes
	pass-through no-pass-through	Yes	No	No	Yes
signal-degrade	ber-threshold-clear <i>value</i>	Yes	Yes	Yes	Yes
	ber-threshold-signal-degrade <i>value</i>	Yes	Yes	Yes	Yes
	interval <i>value</i>	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
tca	odu-tca-bbe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes
	odu-tca-bbe-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
	odu-tca-es (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes
	odu-tca-es-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
	odu-tca-ses (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	odu-tca-ses-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
	odu-tca-uas (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes
	otu-tca-bbe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes
	otu-tca-bbe-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
	otu-tca-es (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	otu-tca-es-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
	otu-tca-fec-ber (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes
	otu-tca-ses (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes
	otu-tca-ses-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
	otu-tca-uas (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	otu-tca-uas-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
<i>transport-monitoring</i>	-	No	No	Yes	Yes
<i>trigger-trigger-identifier</i>	oc-lof (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	oc-lom (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	oc-los (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	oc-tsf (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	oc-wavelength-lock (hold-time (down up) ignore)	Yes	No	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	odu-ais (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	odu-bdi (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	odu-bei (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	odu-iae (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	odu-lck (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	odu-oci (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	odu-sd (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	odu-ttim	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	opu-ptim (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	otu-ais (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	otu-bdi (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	otu-fec-deg (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	otu-fec-exe (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	otu-iae (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
	otu-sd (hold-time (down up) ignore)	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	otu-ttim (hold-time (down up) ignore)	Yes	Yes	Yes	Yes
tti <i>tti-identifier</i>	odu-dapi <i>identifier</i>	Yes	14.1R2 14.2	Yes	Yes
	odu-expected-receive-dapi <i>identifier</i>	Yes	Yes	Yes	Yes
	odu-expected-receive-sapi <i>identifier</i>	Yes	Yes	Yes	Yes
	odu-sapi <i>identifier</i>	Yes	Yes	Yes	Yes
	otu-dapi <i>identifier</i>	Yes	Yes	Yes	Yes
	otu-expected-receive-dapi <i>identifier</i>	Yes	Yes	Yes	Yes
	otu-expected-receive-sapi <i>identifier</i>	Yes	Yes	Yes	Yes

Table 130: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

Statement	Options	P1-PTX-2-100G-WDM (PTX5000 / PTX3000) (13.2R1 / 13.3R1)	P2-100GE-OTN (PTX5000) (14.1R2 / 14.2R1)	P1-PTX-24-10G-W-SFPP (PTX5000) (14.2R1)	PTX10K-LC1104 (PTX10008 and PTX10016) (18.3R1)
	otu-sapi <i>identifier</i>	Yes	Yes	Yes	Yes

SEE ALSO

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

Supported Optics Options on ACX6360 and ACX5448-D Routers

[Table 131 on page 391](#) lists the statements that are supported on ACX6360 and ACX5448-D routers at the [edit interfaces *interface-name* optics-options] hierarchy level.

Table 131: Statements Supported on ACX6360 and ACX5448-D Routers

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1-S1)	Interfaces Supported
<i>fec</i>	sdfec sdfec15	Yes	ot
high-polarization	-	Yes	ot
laser-enable no-laser-enable	-	Yes	ot
los-alarm-threshold	-	Yes	ot
los-warn-threshold	-	Yes	ot

Table 131: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1-S1)	Interfaces Supported
<i>modulation-format</i>	(16qam 8qam qpsk)	Yes	ot
signal-degrade	ber-threshold-clear <i>value</i>	Yes	ot
	ber-threshold-signal-degrade <i>value</i>	Yes	
	interval <i>value</i>	Yes	
	q-threshold-signal-degrade	Yes	
	q-threshold-signal-degrade-clear	Yes	
tca	carrier-frequency-offset-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	ot
	carrier-frequency-offset-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	fec-ber (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	

Table 131: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1-S1)	Interfaces Supported
	fec-corrected-errors-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	fec-uncorrected-words-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	laser-frequency-error-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	laser-frequency-error-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	pam-histogram-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	residual-isi-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	residual-isi-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	

Table 131: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1-S1)	Interfaces Supported
	rx-power-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	rx-power-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	snr-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	tec-current-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	tec-current-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	temperature-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	temperature-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	tx-power-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	

Table 131: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1-S1)	Interfaces Supported
	tx-power-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
<i>tx-power</i>	dbm	Yes	ot
<i>wavelength</i>	nm	Yes	ot

SEE ALSO

Configuring OTN

Supported Optics Options on PTX10008 and PTX10016 Series Routers

[Table 132 on page 395](#) lists the statements that are supported on Juniper Networks® PTX10008 Router and Juniper Networks® PTX10016 Series Router at the [edit interfaces *interface-name* optics-options] hierarchy level.

Table 132: Statements Supported on PTX10008 and PTX10016 Series Routers

Statement	Options	Release (18.3R1)	Interfaces Supported
alarm low-light-alarm	link-down syslog	Yes	ot
tca carrier-frequency-offset-high-tca	(enable-tca no-enable-tca threshold threshold-24hrs)	Yes	ot
<i>tx-power</i>	dbm	Yes	ot

Table 132: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

Statement	Options	Release (18.3R1)	Interfaces Supported
warning low-light-warning	link-down syslog	Yes	ot
laser-enable no-laser-enable	-	Yes	ot
line-loopback no-line-loopback	-	Yes	ot
prbs no-prbs	-	Yes	ot
signal-degrade	ber-threshold-clear <i>value</i>	Yes	ot
	ber-threshold-signal-degrade <i>value</i>	Yes	
	interval <i>value</i>	Yes	
tca	odu-tca-bbe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	ot
	odu-tca-bbe-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	odu-tca-es (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	

Table 132: Statements Supported on PTX10008 and PTX10016 Series Routers (*Continued*)

Statement	Options	Release (18.3R1)	Interfaces Supported
	odu-tca-es-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	odu-tca-ses (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	odu-tca-ses-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	odu-tca-uas (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	otu-tca-bbe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	otu-tca-bbe-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	otu-tca-es (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	otu-tca-es-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	

Table 132: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

Statement	Options	Release (18.3R1)	Interfaces Supported
	otu-tca-ses (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	otu-tca-ses-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	otu-tca-uas (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	otu-tca-uas-fe (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
trigger <i>trigger-identifier</i>	oc-lof (hold-time (down up) ignore)	Yes	ot
	oc-lom (hold-time (down up) ignore)	Yes	
	oc-los (hold-time (down up) ignore)	Yes	
	oc-tsf (hold-time (down up) ignore)	Yes	
	oc-wavelength-lock (hold-time (down up) ignore)	Yes	

Table 132: Statements Supported on PTX10008 and PTX10016 Series Routers (*Continued*)

Statement	Options	Release (18.3R1)	Interfaces Supported
	odu-ais (hold-time (down up) ignore)	Yes	
	odu-bdi (hold-time (down up) ignore)	Yes	
	odu-bei (hold-time (down up) ignore)	Yes	
	odu-iae (hold-time (down up) ignore)	Yes	
	odu-lck (hold-time (down up) ignore)	Yes	
	odu-oci (hold-time (down up) ignore)	Yes	
	odu-sd (hold-time (down up) ignore)	Yes	
	odu-ttim	Yes	
	opu-ptim (hold-time (down up) ignore)	Yes	
	otu-ais (hold-time (down up) ignore)	Yes	
	otu-bdi (hold-time (down up) ignore)	Yes	

Table 132: Statements Supported on PTX10008 and PTX10016 Series Routers (*Continued*)

Statement	Options	Release (18.3R1)	Interfaces Supported
	otu-fec-deg (hold-time (down up) ignore)	Yes	
	otu-fec-exe (hold-time (down up) ignore)	Yes	
	otu-iae (hold-time (down up) ignore)	Yes	
	otu-sd (hold-time (down up) ignore)	Yes	
	otu-ttim (hold-time (down up) ignore)	Yes	

RELATED DOCUMENTATION

Configuring OTN

Forward Error Correction and Bit Error Rate

IN THIS SECTION

- [Overview | 401](#)
- [Supported Forward Error Correction Modes | 405](#)
- [ODU Path Delay Measurement for Performance Monitoring | 407](#)

OTN interfaces use pre-forward error correction (pre-FEC) BER for monitoring the condition of an OTN link. Read this topic to understand more about how OTN links and the supported FEC modes on devices.

Overview

IN THIS SECTION

- [Signal Degrade and Clear Threshold Values for PICs | 404](#)

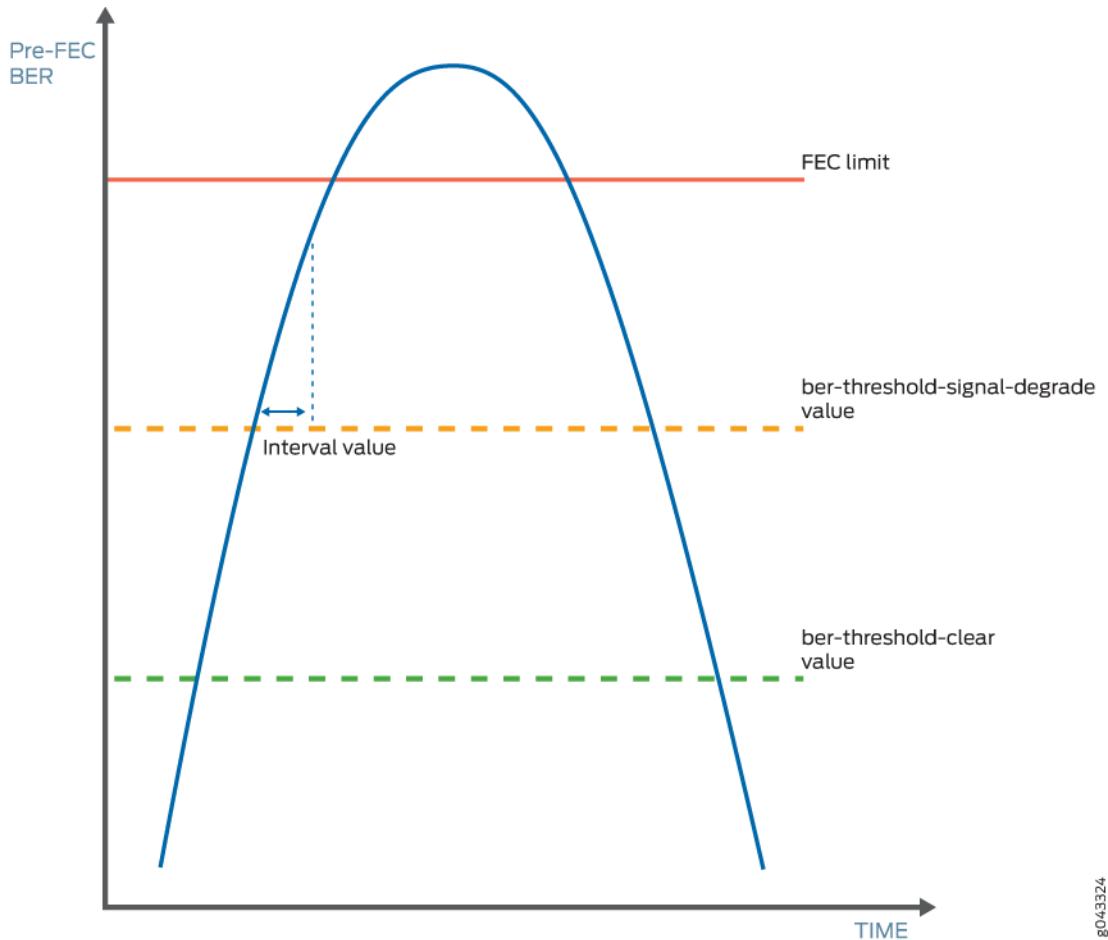
Optical interfaces on PTX Series support monitoring the condition of an optical link by using the pre-forward error correction (pre-FEC) bit error rate (BER). The following PICs support pre-FEC BER monitoring:

- P1-PTX-2-100G-WDM
- P2-100GE-OTN
- P1-PTX-24-10G-W-SFPP

The PICs use forward error correction (FEC) to correct bit errors in the received data. As long as the pre-FEC BER is below the FEC limit, all bit errors are successfully identified and corrected and, therefore, no packet loss occurs. The system monitors the pre-FEC BER on each port. This gives an early warning of link degradation. By configuring an appropriate pre-FEC BER threshold and interval, you enable the PIC to take preemptive action before the FEC limit is reached. If you combine the pre-FEC BER threshold logic with MPLS fast reroute, then you can prevent or minimize packet loss.

You must specify both the signal degradation threshold (*ber-threshold-signal-degrade*) and the interval (*interval*) for the interface. The threshold defines the BER criteria for a signal degrade condition and the interval defines the minimum duration over which the BER must exceed the threshold before an alarm is raised. The relationship between the threshold and the interval is illustrated in [Figure 6 on page 402](#). After an alarm is raised, if the BER returns to a level below the threshold clear value (*ber-threshold-clear*), the alarm is cleared.

Figure 6: Pre-FEC BER Monitoring



8043324

With pre-FEC BER monitoring enabled, when the configured pre-FEC BER signal degrade threshold is reached, the PIC stops forwarding packets to the remote interface and raises an interface alarm. Ingress packets continue to be processed. If pre-FEC BER monitoring is used with MPLS fast reroute or another link protection method, then traffic is rerouted to a different interface.

You can also configure backward fast reroute to insert the local pre-FEC status into transmitted optical frames, notifying the remote interface of signal degradation. The remote interface can use the information to reroute traffic to a different interface. If you use pre-FEC BER monitoring together with backward fast reroute, then notification of signal degradation and rerouting of traffic occurs in less time than that required through a Layer 3 protocol.

Include the `signal-degrade-monitor-enable` and `backward-frr-enable` statements at the `[edit interfaces interface-name otn-options preemptive-fast-reroute]` hierarchy level to enable pre-FEC BER monitoring and backward fast reroute.



NOTE: When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the signal-degrade-monitor-enable and the backward-frr-enable statements.

You can also configure the pre-FEC BER thresholds that raise or clear a signal degrade alarm and the time interval for the thresholds. If the BER thresholds and interval are not configured, the default values are used.

When a received signal degrade alarm is active and backward fast reroute is enabled, a specific flag is inserted into the transmitted optics overhead. The remote PIC at the opposite end of the link monitors the optics overhead, thus enabling both ends to initiate traffic rerouting in the event of a signal degrade condition. When the signal degrade condition is cleared, the overhead flag is returned to a normal state.

The pre-FEC BER signal degrade threshold value defines a specific amount of system margin relative to the BER correction limit (or FEC limit) of the PIC's receive FEC decoder. Each PIC has a set FEC limit—it is intrinsic to the FEC decoder implementation.



NOTE: The examples below use Q²-factor measurements (also known as Q-factor). Q²-factor is expressed in units of decibels relative to a Q²-factor of zero (dBQ). Q²-factor enables you to describe system margin in linear terms in contrast to BER values, which are nonlinear in nature. After you determine the thresholds, you must convert the threshold values from Q²-factor to BER to enter them in the CLI by using scientific notation. BER can be converted to Q²-factor by using the following equation:

$$\text{Q}^2\text{-factor} = 20 * \log10(\sqrt{2} * \text{erfcinv}(2 * \text{BER}))$$



TIP: To convert between Q²-factor and BER in a spreadsheet program, you can approximate the values by using the following formulas:

- To calculate Q²-factor:

$$= 20 * \text{LOG10}(-\text{NORMSINV}(\text{BER}))$$

- To calculate BER:

$$= 1 - \text{NORMSDIST}(10^{(0.05 * \text{Q2-factor})})$$

Include the `ber-threshold-signal-degrade`, `ber-threshold-clear`, and `interval` statements at the `[edit interfaces interface-name otn-options signal-degrade]` hierarchy level to configure the BER thresholds and time interval.



NOTE: Configuring a high BER threshold for signal degradation and a long interval might cause the internal counter register to be saturated. Such a configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

Signal Degrade and Clear Threshold Values for PICs

[Table 133 on page 404](#) shows the relationship between the fixed FEC limit, the configurable signal degrade threshold, and the configurable clear threshold for different PICs. In this example, approximately 1 dBQ of system margin has been set between the FEC limit, signal degrade threshold, and clear threshold.

Table 133: Example—Signal Degrade and Clear Threshold Values at 1 dBQ

PIC	FEC Type	FEC Limit		Signal Degrade Threshold		Clear Threshold	
		Q ² -Factor	BER	Q ² -Factor	BER	Q ² -Factor	BER
P1-PTX-2-100G-WDM	SD-FEC	6.7 dBQ	1.5E-2	7.7 dBQ	7.5E-3	8.7 dBQ	3.0E-3
P2-100GE-OTN	G.709 GFEC	11.5 dBQ	8.0E-5	12.5 dBQ	1.1E-5	13.5 dBQ	1.0E-6
P1-PTX-24-10G-W-SFPP	G.975.1 I.4 (UFEC)	9.1 dBQ	2.2E-3	10.1 dBQ	6.9E-4	11.1 dBQ	1.6E-4
	G.975.1 I.7 (EFEC)	9.6 dBQ	1.3E-3	10.6 dBQ	3.6E-4	11.6 dBQ	7.5E-5
	G.709 GFEC	11.5 dBQ	8.0E-5	12.5 dBQ	1.1E-5	13.5 dBQ	1.0E-6

To adjust the signal degrade threshold, you must first decide on a new system margin target and then calculate the respective BER value (using the equation to convert from Q²-factor to BER). [Table 134 on page 405](#) shows the values if 3 dBQ of system margin relative to the FEC limit is required for the signal degrade threshold (while maintaining the clear threshold at 1 dBQ relative to the signal degrade threshold).

The choice of system margin is subjective, as you might want to optimize your thresholds based on different link characteristics and fault tolerance and stability objectives. For guidance about configuring pre-FEC BER monitoring and BER thresholds, contact your Juniper Networks representative.

Table 134: Example—Signal Degrade and Clear Thresholds after Configuration

PIC	FEC Type	FEC Limit		Signal Degrade Threshold		Clear Threshold	
		Q ² -Factor	BER	Q ² -Factor	BER	Q ² -Factor	BER
P1-PTX-2-100G-WDM	SD-FEC	6.7 dBQ	1.5E-2	9.7 dBQ	1.1E-3	10.7 dBQ	2.9E-4
P2-100GE-OTN	G.709 GFEC	11.5 dBQ	8.0E-5	14.5 dBQ	4.9E-8	15.5 dBQ	1.1E-9
P1-PTX-24-10G-W-SFPP	G.975.1 I.4 (UFEC)	9.1 dBQ	2.2E-3	12.1 dBQ	2.8E-5	13.1 dBQ	3.1E-6
	G.975.1 I.7 (EFEC)	9.6 dBQ	1.3E-3	12.6 dBQ	1.1E-5	13.6 dBQ	9.1E-7
	G.709 GFEC	11.5 dBQ	8.0E-5	14.5 dBQ	4.8E-8	15.5 dBQ	1.1E-9

Supported Forward Error Correction Modes

This section describes FEC modes supported on different routers at the [edit interfaces *interface-name* otn-options] level.

MX Series Routers

Table 135: FEC modes Supported on MX Series Routers

Line Card	FEC Mode	Port Speed
MPC5E-40G10G	(gfec / efec / none / ufec)	10G
MPC5E-100G10G	(gfec / efec / none / ufec)	10G and 100G (GFEC only)

Table 135: FEC modes Supported on MX Series Routers (Continued)

Line Card	FEC Mode	Port Speed
MIC6-10G-OTN	(gfec / efec / none / ufec)	10G
MIC6-100G-CFP2	(gfec / none)	100G (GFEC only)
MIC3-100G-DWDM	gfec / hgfec / sdfec	100G

PTX Series Routers**Table 136: FEC Modes Supported on PTX Series Routers**

Line Card	FEC Mode	Port Speed
P1-PTX-24-10G-W-SFPP	(gfec / efec / none / ufec)	10G
P2-10G-40G-QSFPP	(gfec / efec / none / ufec)	10G
P2-100GE-OTN	(gfec / none)	100G (GFEC only)
P1-PTX-2-100G-WDM	(gfec-sdfec)	100G
PTX-5-100G-WDM	(gfec / sdfec)	100G

Table 137: ACX Series Routers

Device	FEC Mode	Modulation Format	Port Speed
, ACX5448-D	sdfec	QPSK	100G
, ACX5448-D	sdfec15	QPSK	100G
ACX5448-D	hgfec	QPSK	100G

Table 137: ACX Series Routers (*Continued*)

Device	FEC Mode	Modulation Format	Port Speed
, ACX5448-D	<i>sdfec15</i>	8-QAM	200G
, ACX5448-D	<i>sdfec15</i>	16-QAM	200G

SEE ALSO| *fec***ODU Path Delay Measurement for Performance Monitoring****IN THIS SECTION**

- [Overview | 407](#)
- [Enabling ODU Path Delay Measurement | 408](#)
- [Disabling ODU Path Delay Measurement | 410](#)

Read this topic to understand ODU path delay measurement and performance monitoring.

Overview**IN THIS SECTION**

- [Guidelines to Configure Delay Measurement | 408](#)

Performance monitoring is an important requirement in any network, including the optical networks. The key parameters that impact performance are bit error rate (BER) and delay. Delays in data communication over a network impact the network latency. Network latency is the time taken for a packet of data to travel from a designated point to another designated point. If there are less delays, the

network latency is low. You can measure latency by sending a packet and then receiving it as it is returned to you; the time taken for the round-trip indicates the latency.

The optical channel data unit (ODU) path delay measurement offers in-service delay measurement. Delay (or latency) is measured by transmitting a known pattern (delay measurement pattern) in a selected bit of the delay measurement (DM) field and measuring the number of frames that are missed when the delay measurement pattern is received at the transmitting end. For instance, if the transmitted delay measurement bit is 1 1 1 1 1 1 1 1 0 0 and the received delay measurement bit is 1 1 1 0 0 0 0 0 0 0, the delay measurement starts at frame 2 and ends at frame 8. This can be detected by the change in value between the transmitted bit and the received bit.

```
Frame# 10 9 8 7 6 5 4 3 2 1
Tx DM bit 1 1 1 1 1 1 1 1 0 0
Rx DM bit 1 1 1 0 0 0 0 0 0 0
```

The result of the delay measurement is 6 frames (8 - 2).

Guidelines to Configure Delay Measurement

Follow these guidelines to ensure that you obtain accurate delay measurement when you configure in-service delay measurement:

- Unidirectional delay measurement is not supported. The in-service delay measurement is specific to round-trip delay measurement and for optical channel data units only.
- Delay measurement on different framers for the MIC and PIC is different. So, the delay measurement values are different.
- Resiliency is not supported for path delay measurement.
- Links at the local and remote interfaces must be active before you configure delay measurement.
- Do not perform delay measurement tests when ODU maintenance signals are injected.
- Do not configure local loopback and network loopback with remote loopback because the loopback data is overwritten by the delay measurement pattern.
- If a link failure occurs after you begin measuring delay, delay measurement fails. You must enable the delay measurement again on the local interface to measure delay.

Enabling ODU Path Delay Measurement

Delay measurement is disabled by default. This topic explains the broad steps for measuring the optical channel data units (ODU) path delay on optical transport networks (OTN). First, enable remote loopback on the remote interface and commit the configuration. This enables the remote interface to loop back

the delay measurement pattern to the local interface. Then, start delay measurement at the local interface and view the results.



NOTE: Do not enable remote loopback on both ends (local and remote). If you enable remote loopback on both interfaces, the delay measurement pattern is looped back continuously between the two interfaces.

Before you start measuring delay in the ODU path on OTN, complete the following tasks:

- Ensure that the links are active at the local and remote interfaces and alarms are not configured.
- Ensure that there is a delay of 10 seconds before enabling remote loopback. Also, ensure that there is a delay of 10 seconds after enabling remote loopback at the remote interface and before you start measuring delay.
- Ensure that the delay measurement tests are not performed when ODU maintenance signals are injected.
- Ensure that the local loopback and network loopback are also not specified because the looped-back data is overwritten by the delay measurement pattern.



NOTE: If link failure occurs after you begin measuring delay, delay measurement fails. You must re-enable measurement of delay on the local interface to measure delay.

To enable ODU path delay measurement, perform the following steps:

1. Enable remote loopback on the remote interface by including the `remote-loop-enable` statement at the `[edit]` hierarchy level.

```
[edit]
user@host# set interfaces interfacename otn-options odu-delay-management remote-loop-enable
```

2. After enabling remote loopback, commit the configuration.

```
[edit]
user@host# commit
```

3. Start delay measurement on the local interface by including the start-measurement statement at the [edit] hierarchy level.

```
[edit]
user@host# set interface interfacename otn-options odu-delay-management start-measurement
```

4. After enabling measurement of delay on the local interface, commit the configuration.

```
[edit]
user@host# commit
```

5. To view the delay measurement values, from the operational mode, enter the show interfaces extensive command.

```
user@host> show interfaces interfacename extensive
```

```
...
ODU Delay Management:
Start Measurement: True
Remote Loop Enable: False
Result: 0 micro seconds
...
```

Disabling ODU Path Delay Measurement

Delay measurement is disabled by default. If you enabled optical channel data unit (ODU) path delay measurement by using the remote-loop-enable and start-measurement statements, you can use this procedure to disable delay measurement.



NOTE: You can also use the delete or deactivate command to disable remote loopback on the remote interface. For instance, you can use the delete interfaces *interfacename* otn-options odu-delay-management remote-loop-enable or deactivate interface *interfacename* otn-options odu-delay-management remote-loop-enable command to disable remote loopback on the remote interface.

To disable ODU path delay measurement, first disable remote loopback of the delay measurement pattern on the remote interface and then stop delay measurement:

1. Stop delay measurement on the local interface by including the stop-measurement statement at the [edit] hierarchy level.

```
[edit]
user@host# set interface interfacename otn-options odu-delay-management stop-measurement
```

2. After you stop delay measurement on the local interface, commit the configuration.

```
[edit]
user@host# commit
```

3. Disable remote loopback on the remote interface by including the no-remote-loop-enable statement at the [edit] hierarchy level.

```
[edit]
user@host# set interfaces interfacename otn-options odu-delay-management no-remote-loop-enable
```

4. After disabling remote loopback on the remote interface, commit the configuration.

```
[edit]
user@host# commit
```

5. To verify that remote loopback is disabled and delay is not measured, enter the show interfaces extensive command, from the operational mode.

```
user@host> show interfaces interfacename extensive
```

```
...
ODU Delay Management:
Start Measurement: False
Remote Loop Enable: False
Result: 0 micro seconds
...
```

RELATED DOCUMENTATION

[100GbE, 40GbE, and 10GbE | 351](#)

Supported OTN and Optics Options

RELATED DOCUMENTATION

[100GbE, 40GbE, and 10GbE | 351](#)

Configuring OTN

SEE ALSO

Configuring OTN

SEE ALSO

Configuring OTN Interface Options on PTX10K-LC1104

Interface Mapping and Modulation format for PTX10K-LC1104 Line Card

The PTX10K-LC1104 line card supports 3 optical modules and 2 ports per optical modules. 2 *ot* interfaces are created for an optical module. Hence, 6 *ot* interfaces are created for a line card. The optical interface to *et* interface mapping is shown in the following table:

“ot-“ interface	Modulation Format	Mapped “et” interface(s)
ot-0/0/0	QPSK	et-x/0/0
	8QAM	et-x/0/0
		et-x/0/1
ot-0/0/1	16QAM	et-x/0/0
		et-x/0/1
ot-0/0/2	QPSK	et-x/0/2
	8QAM	et-x/0/1
		et-x/0/2
ot-0/0/3	16QAM	et-x/0/2
		et-x/0/3
ot-0/0/2	QPSK	et-x/0/4
	8QAM	et-x/0/4
		et-x/0/5
ot-0/0/3	16QAM	et-x/0/4
		et-x/0/5
ot-0/0/3	QPSK	et-x/0/6

(Continued)

“ot-“ interface	Modulation Format	Mapped “et” interface(s)
ot-0/0/4	8QAM	et-x/0/5
		et-x/0/6
	16QAM	et-x/0/6
		et-x/0/7
	QPSK	et-x/0/8
		et-x/0/8
		et-x/0/9
	16QAM	et-x/0/8
		et-x/0/9
ot-0/0/5	QPSK	et-x/0/10
		et-x/0/9
	8QAM	et-x/0/9
		et-x/0/10
	16QAM	et-x/0/10
		et-x/0/11

SEE ALSO

Configuring OTN

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
18.3R1	Starting in Junos OS Release 18.3R1, Optical transport interfaces on ACX6360 Routers support monitoring the condition of an optical link by using the pre-forward error correction (pre-FEC) bit error rate (BER).
15.1F5	Starting from Junos OS Release 15.1F5, the 100-Gigabit DWDM OTN MIC—MIC3-100G-DWDM—is supported on MPC3E (MX-MPC3E-3D) and MPC3E NG (MPC3E-3D-NG) on the MX240, MX480, MX960, MX2010, and MX2020 routers.
15.1	Starting from Junos OS Release 15.1, you can configure the interfaces on the P2-100GE-OTN PIC on PTX5000 routers, to be a part of the mixed rates and mixed mode aggregated Ethernet bundles.
15.1	Starting from Junos OS Release 15.1, you can configure port-based pseudowire class of service (CoS) classification which includes Layer 3 IPv4, IPv6, and MPLS classification for interfaces with ethernet-ccc encapsulation.

8

PART

Monitoring and Troubleshooting Ethernet Interfaces

- [Monitor Interfaces | 417](#)
- [Troubleshoot Interfaces | 461](#)

CHAPTER 1

Monitor Interfaces

IN THIS CHAPTER

- Passive Monitoring on Ethernet Interfaces | [417](#)
- Enable Passive Monitoring on Ethernet Interfaces | [419](#)
- Link Degrade Monitoring | [421](#)
- Monitor Gigabit Ethernet Interfaces | [422](#)
- Perform Loopback Test for Gigabit Ethernet Interfaces | [437](#)

Passive Monitoring on Ethernet Interfaces

SUMMARY

Learn how to enable passive monitoring of IPv4 traffic using monitoring services I and II PICs on various Ethernet interfaces.

IN THIS SECTION

- What is Passive Monitoring? | [417](#)
- Passive Monitoring Configuration Guidelines | [418](#)
- Platform-Specific DDoS Protection Behavior | [418](#)

What is Passive Monitoring?

Passive monitoring, a type of network monitoring, passively captures traffic from monitoring interfaces.

When you enable passive monitoring, the device accepts and monitors traffic on the interface, then forwards it to monitoring tools such as IDS servers, packet analyzers, or other devices like routers or end hosts. It provides filtering capabilities for monitoring ingress and egress traffic at the Internet point of presence (PoP) where security networks are attached.

IPv4 traffic that supports only IPv4-only PICs, support 10-port Gigabit Ethernet PIC with SFPs, 2-port Gigabit Ethernet PIC with SFPs and 1-port 10-Gigabit Ethernet PIC.

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Review the ["Platform-Specific DDoS Protection Behavior" on page 418](#) section for notes related to your platform.

Passive Monitoring Configuration Guidelines

- You can only configure passive monitoring at the interface level. The system does not support configuration per VLAN or logical interface.
- An aggregated Ethernet (AE) interface cannot function as a passive monitoring interface.
- Monitoring tools or devices must be directly connected to the switch or router.
- The system drops packets that carry more than two MPLS labels and more than two VLAN tags.
- Exception packets such as IP packet options, router alert, and TTL expiry packets are treated as regular traffic.
- We do not support Ethernet encapsulation and Link Aggregation Control Protocol (LACP) on the AE bundle connected to the monitoring tool or device.

Platform-Specific DDoS Protection Behavior

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Use the following table to review platform-specific behaviors for your platform:

Platform	Difference
MX Series	<ul style="list-style-type: none"> • MX240, MX480, and MX960 Series routers that support I3 based, support Type 2 MX FPCs, Type 3 MX FPCs and Gigabit Ethernet Enhanced DPC with SFP (DPCE-R-40GE-SFP). • MX240, MX480, and MX960 Series routers that support Trio-based (10-Gigabit Ethernet MPC with SFP+), support 30-Gigabit Ethernet MPC and 60-Gigabit Ethernet MPC. • MX80 Series routers that support passive monitoring, support 10-Gigabit Ethernet MPC (SFP+) and 30-Gigabit Ethernet MPC interfaces.

Enable Passive Monitoring on Ethernet Interfaces

SUMMARY

Learn how to configure an interface in passive monitoring mode to drop packets destined for the router, stop transmissions, and enable monitoring for IPv4 with specific statements, while using port mirroring for IPv6.

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine drops packets from that interface destined for the router. This mode prevents the Routing Engine from transmitting any packets through that interface. Packets received from the monitored interface can be forwarded to monitoring interfaces. If you add the `passive-monitor-mode` statement in the configuration:

- Gigabit Ethernet interfaces can support both per-port passive monitoring and per-VLAN passive monitoring. The destination MAC filter on the receive port of the Ethernet interfaces is disabled.
- Ethernet encapsulation options are not allowed.
- Ethernet interfaces do not support the `stacked-vlan-tagging` statement for both IPv4 and IPv6 packets in passive monitor mode.

To enable packet flow monitoring on Ethernet interfaces:

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

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

2. Add the `passive-monitor-mode` statement.

```
[edit interfaces interface-name]
user@host# set passive-monitor-mode
```

For IPv4 monitoring services interfaces, enable packet flow monitoring by including the `family` statement at the `[edit interfaces mo-fpc/pic/port unit logical-unit-number]` hierarchy level, specifying the `inet` option:

1. In configuration mode, navigate to the [edit interfaces mo-fpc/pic/port unit logical-unit-number] hierarchy level.

```
[edit]
user@host# edit interfaces mo-fpc/pic/port unit logical-unit-number
```

2. Add the passive-monitor-mode statement.

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number]
user@host# set family inet
```

For conformity with the cflowd record structure, you must add the receive-options-packets and receive-ttl-exceeded statements at the [edit interfaces mo-fpc/pic/port unit *logical-unit-number* family inet] hierarchy level:

1. In configuration mode, navigate to the [edit interfaces mo-fpc/pic/port unit logical-unit-number family inet] hierarchy level.

```
[edit]
user@host# edit interfaces mo-fpc/pic/port unit logical-unit-number family inet
```

2. Add the receive-options-packets and receive-ttl-exceeded statements.

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]
user@host# set receive-options-packets
user@host# set receive-ttl-exceeded
```

IPv6 passive monitoring is not supported on monitoring services PICs. A user must configure port mirroring to forward the packets from the passive monitored ports to other interfaces.

To configure port mirroring, add the port-mirroring statement at the [edit forwarding-options] hierarchy level.

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see [Configuring Multiservice Physical Interface Properties](#) and the [Junos OS Services Interfaces Library for Routing Devices](#).

RELATED DOCUMENTATION

[How Passive Monitoring works on Ethernet Interfaces](#)

[Configuring Multiservice Physical Interface Properties](#)

Link Degrade Monitoring

SUMMARY

Learn how to use Link degrade monitoring to track Ethernet link quality by monitoring bit error rates (BER) and trigger corrective actions when thresholds are exceeded. Configure this feature to prevent packet drops and ensure reliable connections.

Link degrade monitoring tracks physical link quality on Ethernet interfaces and triggers corrective actions when link quality drops below a set level. It is configured using the `link-degrade-monitor` statement at the `[edit interfaces interface-name]` level.

Supported protocols: Link Fault Signaling (LFS) protocol - Layer 2 and Layer 3 protocols.

Limitations: Existing protocols cannot detect extremely low bit error rates (BER) ranging from 10^{-13} through 10^{-5} .

Functionality: The system continuously monitors the link's BER and triggers corrective action when the BER exceeds a user-defined threshold. This threshold is set using the `thresholds` statement at the `[edit interfaces interface-name link-degrade-monitor]` level.

When you configure link monitoring (by using the `link-degrade-monitor` statement), you can set the following options:

Table 138: Configure link monitoring using the `link-degrade-monitor` statement

Configuration	Command
BER Threshold Value: Set the BER threshold value at which corrective action triggers on an interface	<code>thresholds (set clear)</code>

Table 138: Configure link monitoring using the link-degrade-monitor statement (Continued)

Configuration	Command
Link Degrade Interval Value: Set the number of consecutive link degradation events before taking corrective action	<code>thresholds interval intervalvalue</code>
Warning Threshold Value: Set the warning threshold value for link degradation	<code>thresholds warning (set clear)</code>
Action on BER Threshold Level: Configure the action taken when the BER threshold level is reached	<code>thresholds actions media-based</code>
Recovery Type: Choose between automatic or manual recovery	<code>recovery (auto manual)</code>
Recovery Time Interval: Set the time interval after which automatic recovery of the degraded link triggers	<code>recovery timer timer-value</code>

RELATED DOCUMENTATION

[Physical Interface Damping Overview](#)

[Fast Reroute Overview](#)

Monitor Gigabit Ethernet Interfaces

SUMMARY

Learn how to effectively monitor Gigabit Ethernet interfaces by following a comprehensive checklist, understanding key monitoring techniques, and reviewing fiber-optic Ethernet interface specifications.

IN THIS SECTION

- [Checklist to Monitor Gigabit Ethernet Interfaces | 423](#)

- [How to Monitor Gigabit Ethernet Interfaces | 424](#)
- [Fiber-Optic Ethernet Interface Specifications | 435](#)

Checklist to Monitor Gigabit Ethernet Interfaces

IN THIS SECTION

- [Purpose | 423](#)
- [Action | 423](#)
- [Meaning | 424](#)

Purpose

To monitor Gigabit Ethernet interfaces and interface problems.

Action

It provides links and commands for monitoring Gigabit Ethernet interfaces.

Table 139: Monitor Gigabit Ethernet Interfaces

Tasks	Command or Action
"Monitor Gigabit Ethernet Interfaces" on page 422	
"Display the Status of Gigabit Ethernet Interfaces" on page 424	<code>show interfaces terse (fe* ge*)</code>
"Display the Status of a Specific Gigabit Ethernet Interface" on page 424	<code>show interfaces (fe-fpc/pic/port ge-fpc/pic/port)</code>

Table 139: Monitor Gigabit Ethernet Interfaces (*Continued*)

Tasks	Command or Action
"Display Extensive Status for a Specific Gigabit Ethernet Interface" on page 424	<code>show interfaces (fe-fpc/pic/port ge-fpc/pic/port)</code> <code>extensive</code>
"Monitor Statistics for a Gigabit Ethernet Interface" on page 424	<code>monitor interface (fe-fpc/pic/port ge-fpc/pic/port)</code>

Meaning

You can use the commands to monitor and to display the configurations for Gigabit Ethernet interfaces.

How to Monitor Gigabit Ethernet Interfaces

IN THIS SECTION

- [Display the Status of Gigabit Ethernet Interfaces | 424](#)
- [Display the Status of Gigabit Ethernet Interfaces | 426](#)
- [Display the Status of a Specific Gigabit Ethernet Interface | 428](#)
- [Display Extensive Status for a Specific Gigabit Ethernet Interface | 429](#)
- [Monitor Statistics for a Gigabit Ethernet Interface | 434](#)

Monitor the Gigabit Ethernet interfaces to isolate interfaces related problems.

To monitor your Gigabit Ethernet interfaces, follow these steps:

Display the Status of Gigabit Ethernet Interfaces

IN THIS SECTION

- [Purpose | 425](#)
- [Action | 425](#)

- Meaning | 425

Purpose

To display the status of Gigabit Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Action

```
user@host> show interfaces terse (fe* | ge*)
```

Sample Output

```
user@host> show interfaces terse fe*
Interface      Admin Link Proto Local                  Remote
fe-2/1/0       up    up
fe-2/1/0.0     up    up    inet  10.116.115.217/29
fe-3/0/2       up    down
fe-3/0/2.0     up    down
fe-3/0/3       up    up
fe-3/0/3.0     up    up    inet  192.168.223.65/30
fe-4/1/0       down  up
fe-4/1/0.0     up    down  inet  10.150.59.133/30
fe-4/1/1       up    up
fe-4/1/1.0     up    up    inet  10.150.59.129/30
fe-4/1/2       up    down
fe-4/1/2.0     up    down
```

Meaning

The sample output lists only the Gigabit Ethernet interfaces. It shows the status of both the physical and logical interfaces.

Table 140: Gigabit Ethernet Interfaces Status

Physical Interface	Logical Interface	Status Description
fe-2/1/0	fe-2/1/0.0	This interface has both the physical and logical links up and running.
Admin Up	Admin Up	
Link Up	Link Up	
fe-3/0/2	fe-3/0/2.0	This interface has the physical link down, the link layer down, or both down (Link Down). The logical link is also down as a result.
Admin Up	Admin Up	
Link Down	Link Down	
fe-4/1/0	fe-4/1/0.0	This interface is administratively disabled and the physical link is healthy (Link Up), but the logical interface is not established. The logical interface is down because the physical link is disabled.
Admin Down	Admin Up	
Link Up	Link Down	
fe-4/1/2	fe-4/1/2.0	This interface has both the physical and logical links down.
Admin Up	Admin Up	
Link Down	Link Down	

Display the Status of Gigabit Ethernet Interfaces

IN THIS SECTION

- [Purpose | 427](#)
- [Action | 427](#)
- [Meaning | 427](#)

Purpose

To display the status of Gigabit Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Action

```
user@host> show interfaces terse ge*
```

Interface	Admin	Link	Proto	Local	Remote
ge-2/2/0	down	down			
ge-2/2/0.0	up	down	inet	65.113.23.105/30	
ge-2/3/0	up	up			
ge-2/3/0.0	up	up	inet	65.115.56.57/30	
ge-3/1/0	up	up			
ge-3/1/0.0	up	up	inet	65.115.56.193/30	
ge-3/2/0	up	down			

Meaning

This sample output lists only the Gigabit Ethernet interfaces. It shows the status of both the physical and logical interfaces. See Table 3 for a description of what the output means.

Table 141: Status of Gigabit Ethernet Interfaces

Physical Interface	Logical Interface	Status Description
ge-2/2/0	ge-2/2/0.0	This interface is administratively disabled (Admin Down). Both the physical and logical links are down (Link Down).
Admin Down	Admin Up	
Link Down	Link Down	
ge-2/3/0	ge-2/3/0.0	This interface has both the physical and logical links up and running.
Admin Up	Admin Up	
Link Up	Link Up	

Table 141: Status of Gigabit Ethernet Interfaces (Continued)

Physical Interface	Logical Interface	Status Description
ge-3/2/0	ge-3/2/0.0	This interface has both the physical link and the logical interface down.
Admin Up	Admin Up	
Link Down	Link Down	

Display the Status of a Specific Gigabit Ethernet Interface

IN THIS SECTION

- [Purpose | 428](#)
- [Action | 428](#)
- [Meaning | 429](#)

Purpose

To investigate the status of a specific Ethernet interface, use the following CLI command to view its details:

Action

```
user@host> show interfaces (fe-fpc/pic/port | ge-fpc/pic/port)
```

Sample Output 1

The following output is for a Gigabit Ethernet interface with the physical link up:

```
user@host> show interfaces ge-3/1/0
Physical interface: ge-3/1/0, Enabled, Physical link is Up
  Interface index: 41, SNMP ifIndex: 55
  Description: customer connection
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 1000mbps, Loopback: Disabled, Flow control: Enabled
```

```

Device flags  : Present Running
Interface flags: SNMP-Traps
Link flags    : None
Current address: 00:90:69:85:71:99, Hardware address: 00:90:69:85:71:99
Input rate    : 7412216 bps (1614 pps), Output rate: 2431184 bps (1776 pps)
Active alarms  : None
Active defects : None
Logical interface ge-3/1/0.0 (Index 11) (SNMP ifIndex 57)
  Flags: SNMP-Traps, Encapsulation: ENET2
  Protocol inet, MTU: 1500
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.117.65.192/30, Local: 10.115.65.193
  Broadcast: 10.115.65.195

```

Meaning

The first line of sample output 1 and 2 shows that the physical link is up, indicating that it is healthy and can pass packets. Scroll through the output to check for any active alarms or defects. If alarms or defects are present, refer to [Display Extensive Status Information for a Specific Gigabit Ethernet Interface](#), for more detailed diagnostic information about the interface.

Display Extensive Status for a Specific Gigabit Ethernet Interface

IN THIS SECTION

- [Purpose | 429](#)
- [Action | 429](#)
- [Meaning | 431](#)

Purpose

To display extensive status information about a specific Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:

Action

```
user@host> show interfaces (fe-fpc/pic/port | ge-fpc/pic/port) extensive
```

Sample Output

The following sample output is for a Gigabit Ethernet interface:

```

user@router> show interfaces fe-1/3/3 extensive
Physical interface: fe-1/3/3, Enabled, Physical link is Up
  Interface index: 47, SNMP ifIndex: 38
  Description: Test
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags  : Present Running
  Interface flags: SNMP-Traps
  Link flags    : None
  Current address: 00:90:69:8d:2c:de, Hardware address: 00:90:69:8d:2c:de
  Statistics last cleared: 2002-01-11 23:03:09 UTC (1w2d 23:54 ago)
  Traffic statistics:
    Input bytes : 373012658 0 bps
    Output bytes : 153026154 1392 bps
    Input packets: 1362858 0 pps
    Output packets: 1642918 3 pps
  Input errors:
    Errors: 0 , Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 503660
    L3 incompletes: 1 , L2 channel errors: 0 , L2 mismatch timeouts: 0
    FIFO errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Collisions: 0, Drops: 0, Aged packets: 0
    HS link CRC errors: 0, FIFO errors: 0
  Active alarms : None
  Active defects : None
  MAC statistics:           Receive   Transmit
    Total octets      439703575 177452093
    Total packets      1866532 1642916
    Unicast packets    972137 1602563
    Broadcast packets   30 2980
    Multicast packets  894365 37373
    CRC/Align errors    0 0
    FIFO errors        0 0
    MAC control frames 0 0
    MAC pause frames   0 0
    Oversized frames   0
    Jabber frames      0
    Fragment frames    0

```

```

VLAN tagged frames          0
Code violations             0
Filter statistics:
  Input packet count       1866532
  Input packet rejects     0
  Input DA rejects        503674
  Input SA rejects        0
  Output packet count       1642916
  Output packet pad count   0
  Output packet error count 0
CAM destination filters: 5, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete, Link partner status: OK
  Link partner: Full-duplex, Flow control: None
PFE configuration:
  Destination slot: 1, Stream number: 15
  CoS transmit queue bandwidth:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
  CoS weighted round-robin:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
  Description: Test
  Flags: SNMP-Traps, Encapsulation: ENET2
  Protocol inet, MTU: 1500, Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.115.107.192/29, Local: 10.115.107.193
    Broadcast: 10.115.107.199

```

Meaning

The sample output shows where the errors might be occurring and includes autonegotiation information.

Table 142: Errors and their Meaning

Error	Meaning
Policed discards	Discarded frames that were not recognized or were not of interest.

Table 142: Errors and their Meaning (Continued)

Error	Meaning
L2 channel errors	Packets for which the router could not find a valid logical interface. For example, the packet is for a virtual LAN (VLAN) that is not configured on the interface.
MTU	The maximum transmission unit (MTU) must match the interface of either the router at the remote end of the Gigabit Ethernet link, or that of the switch.
Input DA rejects	Number of packets with a destination Media Access Control (MAC) address that is not on the accept list. It is normal to see this number increment.
Input SA rejects	Number of packets with a source MAC address that is not on the accept list. This number only increments when source MAC address filtering is configured.

If the physical link is down, look at the active alarms and defects for the Gigabit Ethernet interface and diagnose the Gigabit Ethernet media accordingly. See ["Locate LINK Alarms and Counters for Gigabit Ethernet" on page 465](#) for an explanation of Gigabit Ethernet alarms.

Table 143: MAC Statistics Errors

Error	Meaning
CRC/Align errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
MAC control frames	The number of MAC control frames.
MAC pause frames	The number of MAC control frames with pause operational code.

Table 143: MAC Statistics Errors (Continued)

Error	Meaning
Jabber frames	<p>The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error.</p> <p>Note that 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 where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.</p>
Fragment frames	<p>The total number of packets received 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.</p> <p>Note that it is entirely normal for fragment frames to increment because both runts (which are normal occurrences due to collisions) and noise hits are counted.</p>

Autonegotiation is the process that connected Ethernet interfaces use to communicate the information necessary to interoperate. Table 6 explains the autonegotiation information of the `show interface interface-name extensive` command output.

Table 144: Autonegotiation Field Information

Autonegotiation Field Information	Explanation
Negotiation status: Incomplete	Ethernet interface has the speed or link mode configured.
Negotiation status: No autonegotiation	Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
Negotiation status: Complete Link partner status: OK	Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process completes successfully.
Link partner: Half-duplex	Depends on the capability of the attached Ethernet device.
Flow control: Symmetric/asymmetric	Displays the types of flow control supported by the remote Ethernet device.

Monitor Statistics for a Gigabit Ethernet Interface

SUMMARY

Learn about how to monitor real-time statistics on Gigabit Ethernet interfaces.

IN THIS SECTION

- [Purpose | 434](#)
- [Action | 434](#)
- [Meaning | 435](#)

Purpose

To monitor statistics for a Gigabit Ethernet interface, use the following Junos OS CLI operational mode command:

Action

```
user@host> monitor interface (fe-fpc/pic/port | ge-fpc/pic/port)
```

We recommend that you use the `monitor interface fe-fpc/pic/port` or `monitor interface ge-fpc/pic/port` command only for diagnostic purposes. Do not leave these commands on during normal router operations because real-time monitoring of traffic consumes additional CPU and memory resources.

Sample Output

The following sample output is for a Gigabit Ethernet interface:

```
user@host> monitor interface fe-2/1/0
Interface: fe-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100mbps
Traffic statistics:                                         Current Delta
  Input bytes:          282556864218 (14208 bps)          [40815]
  Output bytes:         42320313078 (384 bps)           [890]
  Input packets:        739373897 (11 pps)            [145]
  Output packets:       124798688 (1 pps)             [14]
Error statistics:                                         Current Delta
  Input errors:          0                                [0]
  Input drops:            0                                [0]
  Input framing errors:  0                                [0]
  Policed discards:      6625892                         [6]
  L3 incompletes:        75                             [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		
Input MAC/Filter statistics:		
Unicast packets	464751787	[154]
Packet error count	0	[0]

Meaning

Use the information from this command to help narrow down possible causes of an interface problem.

If you are accessing the router from the console connection, make sure you set the CLI terminal type using the `set cli terminal` command.

The second column shows cumulative statistics since the last time you cleared them using the `clear interfaces statistics interface-name` command. The third column shows cumulative statistics since you ran the `monitor interface interface-name` command. If input errors are increasing, follow these steps:

- Check the cabling to the router and ask the carrier to verify the line's integrity. Ensure you are using the correct cables for the interface port—single-mode fiber for a single-mode interface, and multimode fiber for a multimode interface.
- For fiber-optic connections, measure the received light level at the receiver end and ensure it meets the Ethernet interface's specification.
- Measure the transmit light level on the Tx port to confirm it is within the specified range.

Fiber-Optic Ethernet Interface Specifications

[Table 145 on page 435](#) shows the specifications for fiber-optic interfaces.

Table 145: Fiber-Optic Ethernet Interface Specifications

Fiber-Optic Ethernet Interface	Length	Wavelength	Average Launch Power	Receiver Saturation	Receiver Sensitivity
Gigabit Ethernet					

Table 145: Fiber-Optic Ethernet Interface Specifications (*Continued*)

Fiber-Optic Ethernet Interface	Length	Wavelength	Average Launch Power	Receiver Saturation	Receiver Sensitivity
Duplex SC connector					
LH optical interface	49.5-mile 70-km reach on 8.2-micrometer SMF	1480 to 1580 nm	-3 to +2 dBm	-3 dBm	-23 dBm (BER 10 ⁻¹²) for SMF
LX optical interface	6.2-mile 10-km reach on 9/125-micrometer SMF 1804.5-ft 550-m reach on 62.5/125- and 50/125-micrometer MMF	1270 to 1355 nm	-11 to -3 dBm	-3 dBm	-19 dBm
SX optical interface	656-ft 200-m reach on 62.5/125-micrometer MMF 1640-ft 500-m reach on 50/125-micrometer MMF	830 to 860 nm	-9.5 to -4 dBm	-3 dBm	-17 dBm

SEE ALSO

[Fiber-Optic Cable Signal Loss, Attenuation, and Dispersion](#)

[Calculating Power Budget and Power Margin for Fiber-Optic Cables](#)

Perform Loopback Test for Gigabit Ethernet Interfaces

SUMMARY

Learn how to effectively perform loopback testing for Gigabit Ethernet interfaces. This guide provides a comprehensive checklist for diagnosing hardware and circuit problems, creating and verifying loopbacks, configuring ARP entries, clearing interface statistics, and troubleshooting through pings and error checks.

IN THIS SECTION

- [Checklist to Use Loopback Test for Gigabit Ethernet Interfaces | 437](#)
- [Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | 439](#)
- [Create a Loopback | 439](#)
- [Verify That the Gigabit Ethernet Interface Is Up | 444](#)
- [Configure a Static Address Resolution Protocol Table Entry | 448](#)
- [Clear Gigabit Ethernet Interface Statistics | 454](#)
- [Ping the Gigabit Ethernet Interface | 454](#)
- [Check for Gigabit Ethernet Interface Error Statistics | 456](#)
- [Diagnose a Suspected Circuit Problem | 459](#)

Checklist to Use Loopback Test for Gigabit Ethernet Interfaces

IN THIS SECTION

- [Purpose | 437](#)
- [Action | 437](#)

Purpose

To use loopback testing to isolate Gigabit Ethernet Interface problems.

Action

The table below provides links and commands for using loopback testing for Gigabit Ethernet Interfaces.

Table 146: Checklist to Use Loopback Test for Gigabit Ethernet Interfaces

Tasks	Command or Action
"Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface" on page 439	
"Create a Loopback" on page 439	
<ul style="list-style-type: none"> • "Create a Physical Loopback for a Fiber-Optic Interface" on page 439 	Connect the transmit port to the receive port.
<ul style="list-style-type: none"> • "Create a Loopback Plug for an RJ-45 Ethernet Interface" on page 439 	Cross pin 1 (TX+) and pin 3 (RX+) together, and pin 2 (TX-) and pin 6 (RX-) together.
<ul style="list-style-type: none"> • "Configure a Local Loopback" on page 439 	<pre>[edit interfaces <i>interface-name</i> (gigether-options)] set loopback</pre>
"Verify That the Gigabit Ethernet Interface Is Up" on page 444	<pre>show interfaces (fe-fpc/pic/port ge-fpc/pic/port)</pre>
"Configure a Static Address Resolution Protocol Table Entry" on page 448	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>] set arp <i>ip-address</i> mac <i>mac-address</i></pre>
"Clear Gigabit Ethernet Interface Statistics" on page 454	<pre>clear interfaces statistics fe-fpc/pic/port / ge-fpc/pic/port</pre>
"Ping the Gigabit Ethernet Interface" on page 454	<pre>ping <i>remote-IP-address</i> bypass-routing interface (fe-fpc/pic/port ge-fpc/pic/port) count 100 rapid</pre>
"Check for Gigabit Ethernet Interface Error Statistics" on page 456	<pre>show interfaces statistics (fe-fpc/pic/port ge-fpc/pic/port) extensive</pre>

Table 146: Checklist to Use Loopback Test for Gigabit Ethernet Interfaces (Continued)

Tasks	Command or Action
"Diagnose a Suspected Circuit Problem" on page 459	Perform Steps 2 through 8 from "Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface" on page 439 .

Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface

IN THIS SECTION

- [Problem | 439](#)
- [Solution | 439](#)

Problem

When you suspect a hardware problem, take the following steps to help verify if there is a problem.

Solution

To diagnose a suspected hardware problem with the Ethernet interface, follow these steps:

- Create a Loopback
- Verify That the Gigabit Ethernet Interface Is Up
- Configure a Static Address Resolution Protocol Table Entry
- Clear Gigabit Ethernet Interface Statistics
- Check for Gigabit Ethernet Interface Error Statistics

Create a Loopback

IN THIS SECTION

- [Create a Physical Loopback for a Fiber-Optic Interface | 440](#)

- [Create a Loopback Plug for an RJ-45 Ethernet Interface | 441](#)
- [Configure a Local Loopback | 442](#)

You can create a physical loopback or configure a local loopback to help diagnose a suspected hardware problem. Creating a physical loopback is recommended because it allows you to test and verify the transmit and receive ports. If a field engineer is not available to create the physical loopback, you can configure a local loopback for the interface. The local loopback creates a loopback internally in the Physical Interface Card (PIC).

Create a Physical Loopback for a Fiber-Optic Interface

IN THIS SECTION

- [Purpose | 440](#)
- [Action | 440](#)
- [Meaning | 440](#)

Purpose

To verify the functionality of the fiber-optic interface by testing both the transmit and receive ports. This test is typically performed by a field engineer to confirm the integrity of the connection.

Action

To create a physical loopback at the port, connect the transmit port to the receive port using a known good fiber cable.

Use single-mode fiber for a single-mode port and multimode fiber for a multimode port.

Meaning

When you create and test a physical loopback, you are effectively testing the transmit and receive capabilities of the PIC. This procedure provides a comprehensive test of the PIC, confirming that both ports are operational and the fiber-optic connection is functioning as expected.

SEE ALSO

[Perform Loopback Test for Gigabit Ethernet Interfaces | 437](#)

Create a Loopback Plug for an RJ-45 Ethernet Interface**IN THIS SECTION**

- [Purpose | 441](#)
- [Action | 441](#)
- [Meaning | 442](#)

Purpose

To verify the functionality of an RJ-45 Ethernet interface by creating a physical loopback plug that allows for the testing of the interface's transmit and receive capabilities.

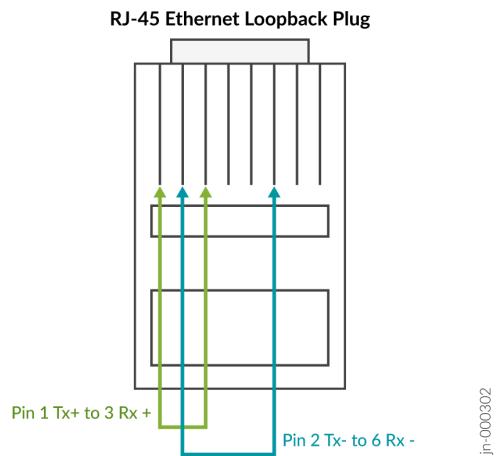
Action

To create a loopback plug, cross pin 1 (TX+) and pin 3 (RX+) together, and cross pin 2 (TX-) and pin 6 (RX-) together. You need the following equipment to create the loopback:

- A 6-inch long CAT5 cable
- An RJ-45 connector
- A crimping tool

Figure 1 illustrates how to create a loopback plug for an RJ-45 Ethernet interface.

Figure 7: RJ-45 Ethernet Loopback Plug



Meaning

When you create and then test a physical loopback, you are testing the RJ-45 interface of the PIC. This action is recommended if a field engineer is available to create the physical loop as it provides a more complete test of the PIC.

SEE ALSO

[Perform Loopback Test for Gigabit Ethernet Interfaces | 437](#)

Configure a Local Loopback

IN THIS SECTION

- [Purpose | 443](#)
- [Action | 443](#)
- [Meaning | 444](#)

Purpose

To configure a local loopback on an Ethernet interface for testing the internal functionality of the PIC without a physical connection between the transmit and receive ports.

Action

To configure a local loopback without physically connecting the transmit port to the receive port, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces interface-name (gigether-options)
```

2. Configure the local loopback:

```
[edit interfaces interface-name (gigether-options)]
user@host# set loopback
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit interfaces fe-1/0/0 gigether-options]
user@host# show loopback
```

4. Commit the change:

```
user@host# commit
```

For example:

```
[edit interfaces fe-1/0/0 gigether-options]
user@host# commit
commit complete
```

When you create a local loopback, you create an internal loop on the interface being tested. A local loopback loops the traffic internally on that PIC. A local loopback tests the interconnection of the PIC but does not test the transmit and receive ports. On an Ethernet interface, you cannot create a remote loopback, therefore there is no option to use a local or remote statement. Simply including the loopback statement at the [edit interfaces *interface-name* (gigether-options] hierarchy level, places the interface into local loopback mode.

Delete the loopback statement after completing the test.

Meaning

Configuring a local loopback internally loops traffic on the interface, verifying the PIC's internal connections. The show command will display the loopback statement in the interface configuration, confirming that the interface is in loopback mode. Remember to delete the loopback configuration after testing.

Verify That the Gigabit Ethernet Interface Is Up

IN THIS SECTION

- [Purpose | 444](#)
- [Action | 445](#)
- [Meaning | 448](#)

Purpose

Display the status of the Gigabit Ethernet Interface to provide the information you need to determine whether the physical link is up or down.

Action

To verify that the status of the Gigabit Ethernet Interface is up, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show interfaces (fe-fpc/port|ge-fpc/pic/port)
```

Sample Output

```
user@host# show interfaces ge-4/0/6 extensive
Physical interface: ge-4/0/6, Enabled, Physical link is Up  Interface index: 144, SNMP ifIndex: 516,
Generation: 147
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags : Present Running Loop-Detected
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 0
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
Last flapped : 2015-01-20 23:40:04 PST (00:02:12 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
Dropped traffic statistics due to STP State:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
```

```

L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0

Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0

Egress queues: 8 supported, 4 in use

Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort          0                  0                  0
  1 expedited-fo          0                  0                  0
  2 assured-forw          0                  0                  0
  3 network-cont          0                  0                  0

Queue number:      Mapped forwarding classes
  0                  best-effort
  1                  expedited-forwarding
  2                  assured-forwarding
  3                  network-control

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

Filter statistics:
  Input packet count      0
  Input packet rejects     0
  Input DA rejects          0
  Input SA rejects          0
  Output packet count       0
  Output packet pad count     0
  Output packet error count    0

CAM destination filters: 0, CAM source filters: 0

Autonegotiation information:

```

```

Negotiation status: Complete
Link partner:
  Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK
Local resolution:
  Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 4
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth          Buffer Priority  Limit
                           %                 bps            %           usec
  0 best-effort          95      950000000  95           0      low   none
  3 network-control       5      500000000  5            0      low   none
Interface transmit statistics: Disabled

```

Meaning

The sample output shows that the link is up and there are no alarms in this loopback configuration. When an internal loopback is configured, the physical loopback should come up without an alarm.

Sample Output

When you see that the physical link is down, there may be a problem with the port. The following output is an example of the `show interfaces fe-fpc/pic/port` command when the physical link is down:

```

user@router> show interfaces fe-1/3/0
Physical interface: fe-1/3/0, Enabled, Physical link is Down
  Interface index: 44, SNMP ifIndex: 35
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags  : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps
  Link flags    : None
  Current address: 00:90:69:8d:2c:db, Hardware address: 00:90:69:8d:2c:db
  Input rate     : 0 bps (0 pps), Output rate: 0 bps (0 pps)
  Active alarms : LINK
  Active defects : LINK
  MAC statistics:
    Input octets: 0, Input packets: 0, Output octets: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Autonegotiation information:

```

```

Negotiation status: Incomplete, Link partner status: Down
Reason: Link partner autonegotiation failure
Link partner: Half-duplex, Flow control: None

```

Meaning

The sample output shows that the physical link is down and there are active alarms and defects.

[Table 147 on page 448](#) presents problem situations and actions for a physical link that is down.

Table 147: Problems and Solutions for a Physical Link That Is Down

Problem	Action
Cable mismatch	Verify that the fiber connection is correct.
Damaged and/or dirty cable	Verify that the fiber can successfully loop a known good port of the same type.
Too much or too little optical attenuation	Verify that the attenuation is correct per the PIC optical specifications.
The transmit port is not transmitting within the dBm optical range per the specifications	Verify that the Tx power of the optics is within range of the PIC optical specification.
Mismatch between the cable type and the port	Verify that a single-mode fiber cable is connected to a single-mode interface and that a multimode fiber cable is connected to a multimode interface. (This problem does not always cause the physical link to go down; errors and dropped packets are sometimes the result.)

Configure a Static Address Resolution Protocol Table Entry

IN THIS SECTION

- [Purpose | 449](#)
- [Action | 449](#)
- [Meaning | 453](#)

Purpose

Configure a static Address Resolution Protocol (ARP) entry to allow a packet to be sent out of a looped Ethernet interface.



NOTE: Remove the static ARP entry at the end of the loop test after you have completed the tests and monitored interface traffic.

Action

To configure a static ARP table entry for a Gigabit Ethernet interface, follow these steps. You can follow the same procedure to configure a static ARP entry for a Gigabit Ethernet interface.

1. Find the Media Access Control (MAC) address for the Gigabit Ethernet interface:

```
user@host> show interfaces ge-fpc/pic/port
```

```
Physical interface: ge-4/0/6, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 516, Generation: 147
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags    : Present Running Loop-Detected
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
  Last flapped    : 2015-01-20 23:40:04 PST (00:13:49 ago)
  Statistics last cleared: 2015-01-20 23:46:15 PST (00:07:38 ago)
  Traffic statistics:
    Input bytes : 125500 0 bps
    Output bytes : 125482 0 bps
    Input packets: 1281 0 pps
    Output packets: 1281 0 pps
  IPv6 transit statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Dropped traffic statistics due to STP State:
```

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
  0 best-effort 1260 1260 0
  1 expedited-fo 0 0 0
  2 assured-forw 0 0 0
  3 network-cont 0 0 0
Queue number: Mapped forwarding classes
  0 best-effort
  1 expedited-forwarding
  2 assured-forwarding
  3 network-control
Active alarms : None
Active defects : None
MAC statistics: Receive Transmit
  Total octets 130624 130624
  Total packets 1281 1281
  Unicast packets 1280 1280
  Broadcast packets 1 1
  Multicast packets 0 0
  CRC/Align errors 0 0
  FIFO errors 0 0
  MAC control frames 0 0
  MAC pause frames 0 0
  Oversized frames 0
  Jabber frames 0
  Fragment frames 0
  VLAN tagged frames 0
  Code violations 0
Filter statistics:
  Input packet count 1281
  Input packet rejects 0
  Input DA rejects 0

```

```

Input SA rejects          0
Output packet count      1281
Output packet pad count  0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 4
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth          Buffer Priority  Limit
                           %                 bps      %           usec
  0 best-effort           95     9500000000  95           0     low    none
  3 network-control       5      50000000    5           0     low    none
Interface transmit statistics: Disabled

Logical interface ge-4/0/6.0 (Index 72) (SNMP ifIndex 573) (Generation 137)
Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 125500
  Output bytes : 123480
  Input packets: 1281
  Output packets: 1260
Local statistics:
  Input bytes : 60
  Output bytes : 2002
  Input packets: 1
  Output packets: 21
Transit statistics:
  Input bytes : 0          0 bps
  Output bytes : 0          0 bps
  Input packets: 0          0 pps
  Output packets: 0          0 pps
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp ospf pgm
pim rip router-discovery rsvp sap vrrp
Flow Statistics :
Flow Input statistics :

```

```

Self packets : 0
ICMP packets : 40
VPN packets : 0
Multicast packets : 0
Bytes permitted by policy : 107520
Connections established : 20
Flow Output statistics:
  Multicast packets : 0
  Bytes permitted by policy : 107520
Flow error statistics (Packets dropped due to):
  Address spoofing: 0
  Authentication failed: 0
  Incoming NAT errors: 0
  Invalid zone received packet: 0
  Multiple user authentications: 0
  Multiple incoming NAT: 0
  No parent for a gate: 0
  No one interested in self packets: 0
  No minor session: 0
  No more sessions: 0
  No NAT gate: 0
  No route present: 11
  No SA for incoming SPI: 0
  No tunnel found: 0
  No session for a gate: 0
  No zone or NULL zone binding 0
  Policy denied: 0
  Security association not active: 0
  TCP sequence number out of window: 0
  Syn-attack protection: 0
  User authentication errors: 0
Protocol inet, MTU: 1500, Generation: 158, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.108.120.0/30, Local: 10.108.120.1, Broadcast: 10.108.120.3,
    Generation: 158
Protocol multiservice, MTU: Unlimited, Generation: 159, Route table: 0
  Policer: Input: __default_arp_policer__

```

2. In configuration mode, go to the following hierarchy level:

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

3. Configure the static ARP entry:

```
user@host# set arp ip-address mac mac-address
```

4. Commit the configuration:

```
user@host# commit
```

5. Verify that the static ARP entry is installed:

```
[edit interfaces ge-4/0/6 unit 0 family inet address 10.108.120.1/30]
user@host# run show arp no-resolve
```

MAC Address	Address	Interface	Flags
00:1f:12:fe:c5:2e	10.108.120.2	ge-4/0/6.0	permanent
52:54:00:7d:33:4c	10.204.128.35	fxp0.0	none
52:54:00:65:11:50	10.204.128.36	fxp0.0	none
52:54:00:da:30:82	10.204.128.37	fxp0.0	none
52:54:00:3a:cf:4b	10.204.128.38	fxp0.0	none
52:54:00:de:88:5f	10.204.128.45	fxp0.0	none
52:54:00:48:03:b7	10.204.128.46	fxp0.0	none

Meaning

The sample output is for Step 1 through Step 5 and shows that a static ARP entry was configured on Gigabit Ethernet interface ge-4/0/6.

Clear Gigabit Ethernet Interface Statistics

IN THIS SECTION

- [Purpose | 454](#)
- [Action | 454](#)
- [Meaning | 454](#)

Purpose

You can reset the Gigabit Ethernet Interface statistics. Resetting the statistics provides a clean start so that previous input/output errors and packet statistics do not interfere with the current diagnostics.

Action

To clear all statistics for the interface, use the following Junos OS CLI operational mode command:

```
user@host> clear interfaces statistics (fe-fpc/pic/port | ge-fpc/pic/port)
```

Sample Output

```
user@host> clear interfaces statistics ge-4/0/6
user@host>
```

Meaning

This command clears the interface statistics counters for the Gigabit Ethernet Interface only.

Ping the Gigabit Ethernet Interface

IN THIS SECTION

- [Purpose | 455](#)
- [Action | 455](#)
- [Meaning | 455](#)

Purpose

Use the ping command to verify the loopback connection.

Action

To send ping packets from the Ethernet interface, use the following Junos OS CLI operational mode command:

```
user@host> ping remote-IP-address bypass-routing interface (fe-fpc/pic/port | ge-fpc/pic/port)
count 100 rapid
```

Sample Output

```
user@router> ping 10.108.120.2 bypass-routing interface ge-7/2/1 count 100 rapid
PING 10.108.120.2 (10.108.120.2): 56 data bytes
36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
4 5 00 0054 e871 0 0000 01 01 cc5c 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
4 5 00 0054 e874 0 0000 01 01 cc59 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
4 5 00 0054 e878 0 0000 01 01 cc55 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
4 5 00 0054 e87c 0 0000 01 01 cc51 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
4 5 00 0054 e880 0 0000 01 01 cc4d 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
4 5 00 0054 e884 0 0000 01 01 cc49 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
```

Meaning

- The sample output shows that the time to live (TTL) expired, indicating that the link is receiving frames from the ping test.

- The MAC address matches the physical address of the port being tested, which allows the port to accept frames from the ping test.
- As the packet loops over the link, a "TTL exceeded" message is expected for each ping sent.
- These messages are generated because the ping packets repeatedly loop between the router and the physical loopback.
- When the packet reaches the other end of the link (which doesn't exist), the loopback sends it back to the same interface, where the Packet Forwarding Engine fabric routes it again.
- After each route lookup, the TTL decreases, and the packet is sent out through the looped interface until the TTL expires, triggering the "TTL expired" message.
- If any errors occur, the packet is discarded, and a time-out error is displayed instead of the TTL expired message.
- By default, the TTL for ICMP echo packets in Junos OS is 64, meaning the packet must loop 63 times before a "TTL expired" message is generated; however, you can adjust the TTL value to change the test's tolerance for loss, such as setting it to 255, which requires the packet to loop 254 times without error before a "TTL expired" message occurs.

Check for Gigabit Ethernet Interface Error Statistics

IN THIS SECTION

- [Purpose | 456](#)
- [Action | 457](#)
- [Meaning | 459](#)

Purpose

Persistent interface error statistics indicate that you need to open a case with the Juniper Networks Technical Assistance Center (JTAC).

Action

To check the local interface for error statistics, use the following Junos OS CLI operational mode command:

```
user@host>show interfaces (fe-fpc/pic/port | ge-fpc/pic/port) extensive
```

Sample Output

```
user@router> show interfaces ge-4/0/6 extensive
Physical interface: ge-4/0/6, Enabled, Physical link is Up  Interface index: 144, SNMP ifIndex: 516,
Generation: 147
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags  : Present Running Loop-Detected
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags    : None
  CoS queues    : 8 supported, 4 maximum usable queues
  Schedulers    : 0
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
  Last flapped   : 2015-01-20 23:40:04 PST (00:02:12 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
  Dropped traffic statistics due to STP State:
    Input bytes   : 0
    Output bytes  : 0
    Input packets: 0
    Output packets: 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
```

L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
 Resource errors: 0

Output errors:
 Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0

Egress queues: 8 supported, 4 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

Queue number: Mapped forwarding classes
 0 best-effort
 1 expedited-forwarding
 2 assured-forwarding
 3 network-control

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	

Filter statistics:

Input packet count	0
Input packet rejects	0
Input DA rejects	0
Input SA rejects	0
Output packet count	0
Output packet pad count	0
Output packet error count	0

CAM destination filters: 0, CAM source filters: 0

Autonegotiation information:

```

Negotiation status: Complete
Link partner:
  Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK
Local resolution:
  Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 4
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth          Buffer Priority  Limit
                           %                 bps            %           usec
  0 best-effort          95      950000000  95           0      low   none
  3 network-control       5      500000000  5            0      low   none
Interface transmit statistics: Disabled

```

Meaning

Check for any error statistics. There should not be any input or output errors. If there are any persistent input or output errors, open a case with the Juniper Networks Technical Assistance Center (JTAC) at support@juniper.net, or at 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).

Diagnose a Suspected Circuit Problem

IN THIS SECTION

- [Purpose | 459](#)
- [Action | 460](#)

Purpose

When you suspect a circuit problem, it is important to work with the transport-layer engineer to resolve the problem. The transport-layer engineer may create a loop to the router from various points in the network. You can then perform tests to verify the connection from the router to that loopback in the network.

Action

After the transport-layer engineer has created the loop to the router from the network, you must verify the connection from the router to the loopback in the network. Follow Step 2 through Step 8 in Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface. Any problems encountered in the test indicate a problem with the connection from the router to the loopback in the network.

By performing tests to loopbacks at various points in the network, you can isolate the source of the problem.

CHAPTER 2

Troubleshoot Interfaces

IN THIS CHAPTER

- Configure Interface Diagnostics Tools to Test the Physical Layer Connections | [461](#)
- Locate LINK Alarms and Counters for Gigabit Ethernet | [465](#)
- Troubleshoot: 10-Gigabit Ethernet Port Stuck in Down State | [473](#)
- Verify Link and Transceivers with Pseudo Random Binary Sequence (PRBS) Test | [475](#)

Configure Interface Diagnostics Tools to Test the Physical Layer Connections

SUMMARY

Learn how to configure interface diagnostics tools for physical layer testing, including loopback and BERT tests.

IN THIS SECTION

- Configure Loopback Testing for Interfaces | [461](#)
- Set Up BERT Test for Physical Connections | [463](#)
- Start and Stop a BERT Test | [464](#)

Configure Loopback Testing for Interfaces

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: Aggregated Ethernet, and Gigabit Ethernet.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

- Configure a line loopback on the local router.
 - The line loopback sends the signal back to the originating router instead of transmitting it to the far-end device.
 - If the originating router receives its own Data Link Layer packets, the problem lies beyond the local router.
- Configure a line loopback farther from the local router - If the originating router does not receive its own Data Link Layer packets, the issue is likely on one of the segments between the local router and the remote router's interface card.
- Narrow down the problem by configuring a line loopback closer to the local router - This helps identify the specific segment causing the issue.

The following types of loopback testing are supported by Junos OS:

- DCE local—Loops packets back on the local data circuit-terminating equipment (DCE).
- DCE remote—Loops packets back on the remote DCE.
- Local—Useful for troubleshooting physical PIC errors. To test a local loopback, issue the `show interfaces interface-name` command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- Payload—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- Remote—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner.

[Table 148 on page 462](#) shows the loopback modes supported on the various interface types.

Table 148: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Gigabit Ethernet	Local	Configuring Ethernet Loopback Capability
Serial (V.35 and X.21)	Local and remote	Configuring Serial Loopback Capability

Table 148: Loopback Modes by Interface Type (Continued)

Interface	Loopback Modes	Usage Guidelines
Serial (EIA-530)	DCE local, DCE remote, local, and remote	Configuring Serial Loopback Capability

To configure loopback testing, include the `loopback mode;`

```
user@host# loopback mode;
```

You can include this statement at the following hierarchy levels:

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

Set Up BERT Test for Physical Connections

To configure BERT:

1. Configure the duration of the test.

```
[edit interfaces interface-name interface-type-options]
user@host#bert-period seconds;
```

You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs. By default, the BERT period is 10 seconds.

2. Configure the error rate to monitor when the inbound pattern is received.

```
[edit interfaces interface-name interface-type-options]
user@host#bert-error-rate rate;
```

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

3. Configure the bit pattern to send on the transmit path.

```
[edit interfaces interface-name interface-type-options]
user@host#bert-algorithm algorithm;
```

algorithm is the pattern to send in the bit stream. For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.152 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
...
...
```

For specific hierarchy information, see the individual interface types.

Start and Stop a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the **disable** statement at the [edit interfaces *interface-name*] hierarchy level:

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

After you configure the BERT properties and commit the configuration, begin the test by issuing the **test interface *interface-name* *interface-type-bert-start*** operational mode command:

```
user@host> test interface interface-name interface-type-bert-start
```

The test runs for the duration you specify with the **bert-period** statement. If you want to terminate the test sooner, issue the **test interface *interface-name* *interface-type-bert-stop*** command:

```
user@host> test interface interface-name interface-type-bert-stop
```

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the `show interfaces extensive | find BERT` command:

```
user@host> show interfaces interface-name extensive | find BERT
```

For more information about running and evaluating the results of the BERT procedure, see the [CLI Explorer](#).

To exchange BERT patterns between a local router and a remote router, include the `loopback remote` statement in the interface configuration at the remote end of the link. From the local router, issue the `test interface` command.

RELATED DOCUMENTATION

<https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/command/show-interfaces-diagnostics-optics-10-gigabit-ethernet.html>

Locate LINK Alarms and Counters for Gigabit Ethernet

SUMMARY

Learn how to locate, display LINK alarms and counters for Gigabit Ethernet interfaces.

IN THIS SECTION

- [Checklist: Find LINK Alarms and Counters for Ethernet | 466](#)
- [Display LINK Alarms on Gigabit Ethernet | 466](#)
- [Gigabit Ethernet Counters | 469](#)

Checklist: Find LINK Alarms and Counters for Ethernet

IN THIS SECTION

- [Purpose | 466](#)
- [Action | 466](#)

Purpose

To find LINK alarm and major counters associated with Gigabit Ethernet interfaces.

Action

[Table 149 on page 466](#) provides links and commands for locating LINK alarm and major counters for Gigabit Ethernet interfaces.

Table 149: Checklist to Locate Gigabit Ethernet Alarms and Counters

Tasks	Command or Action
"Display LINK Alarms on Gigabit Ethernet" on page 466	show interfaces (fe-fpc/pic/port ge-fpc/pic/port) extensive
"Gigabit Ethernet Counters" on page 469	It displays the major counters that appear in the output for the show interfaces fe-fpc/pic/port extensive and the show interfaces ge-fpc/pic/port extensive commands. These counters generally increment when there is a problem with a Gigabit Ethernet interface.

Display LINK Alarms on Gigabit Ethernet

IN THIS SECTION

- [Problem | 467](#)
- [Cause | 467](#)

- [Solution | 467](#)

Problem

To display the Gigabit Ethernet LINK alarm, use the following Junos OS command-line interface (CLI) operational mode command:

Cause

A LINK alarm on Gigabit Ethernet interfaces typically indicates a physical layer issue. This could be due to reasons, such as a disconnected cable, faulty hardware, or a mismatch in speed and duplex settings between connected devices. Identifying the cause of the LINK alarm is crucial to maintain network stability and to ensure communication across the network.

Solution

```
user@host> show interfaces (fe-fpc/pic/port / ge-fpc/pic/port)
extensive
```

Sample Output

The following sample output is for a Gigabit Ethernet interface:

```
user@host> show interfaces fe-1/3/3 extensive
Physical interface: fe-1/3/3, Enabled, Physical link is Down
  Interface index: 47, SNMP ifIndex: 38
  Description: Test
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
  Device flags  : Present Running
  Interface flags: SNMP-Traps
  Link flags    : None
  Current address: 00:90:69:8d:2c:de, Hardware address: 00:90:69:8d:2c:de
  Statistics last cleared: 2002-01-11 23:03:09 UTC (1w2d 23:54 ago)
  Traffic statistics:
    Input bytes   : 373012658          0 bps
    Output bytes  : 153026154          1392 bps
    Input packets: 1362858            0 pps
```

```

Output packets: 1642918 3 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 503660
  L3 incompletes: 1, L2 channel errors: 0, L2 mismatch timeouts: 0
  FIFO errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Collisions: 0, Drops: 0, Aged packets: 0
  HS link CRC errors: 0, FIFO errors: 0
Active alarms : LINK
Active defects : LINK
MAC statistics:           Receive   Transmit
  Total octets           439703575 177452093
  Total packets          1866532   1642916
  Unicast packets        972137   1602563
  Broadcast packets      30        2980
  Multicast packets     894365   37373
  CRC/Align errors       0         0
  FIFO errors            0         0
  MAC control frames    0         0
  MAC pause frames      0         0
  Oversized frames       0
  Jabber frames          0
  Fragment frames        0
  VLAN tagged frames     0
  Code violations         0
Filter statistics:
  Input packet count    1866532
  Input packet rejects   0
  Input DA rejects       503674
  Input SA rejects       0
  Output packet count    1642916
  Output packet pad count 0
  Output packet error count 0
CAM destination filters: 5, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete, Link partner status: OK
  Link partner: Full-duplex, Flow control: None
PFE configuration:
  Destination slot: 1, Stream number: 15
  CoS transmit queue bandwidth:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
  CoS weighted round-robin:
    Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5

```

```

Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
  Description: Test
  Flags: SNMP-Traps, Encapsulation: ENET2
  Protocol inet, MTU: 1500, Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.115.107.192/29, Local: 10.115.107.193
  Broadcast: 10.115.107.199

```

Meaning

The sample output shows where the alarm and other errors might be occurring and any counters that are incrementing. The only alarm associated with Gigabit Ethernet interfaces is the LINK alarm. A LINK alarm indicates a physical problem. To isolate where the physical problem might be occurring, conduct loopback testing. See ["Checklist to Use Loopback Test for Gigabit Ethernet Interfaces" on page 437](#) for information on conducting a loopback test.

Because link status is polled once every second, certain features or mechanisms that require fast link down detection, such as Multiprotocol Label Switching (MPLS) fast reroute, take longer to execute.

Gigabit Ethernet Counters

IN THIS SECTION

- [Problem | 469](#)
- [Cause | 469](#)
- [Solution | 470](#)

Problem

[Table 150 on page 470](#) shows the major counters that appear in the output for the `show interfaces fe-fpcl/pcl/port extensive` and the `show interfaces ge-fpcl/pcl/port extensive` commands. These counters generally increment when there is a problem with a Gigabit Ethernet interface. In the `Counters` column, the counters are listed in the order in which they are displayed in the output.

Cause

The counters typically increment due to issues like hardware faults, physical connection problems, incorrect configurations, or excessive network traffic. These increments indicate underlying problems with the Ethernet interface that need further investigation.

Solution

Table 150: Major Gigabit Ethernet Counters

Counter	Description	Reason for Increment
Input Errors:		
Errors	The sum of the incoming frame terminates and frame check sequence (FCS) errors.	The incoming errors counter increases when there are issues with the integrity of received frames, such as incomplete frame terminations or frame check sequence (FCS) errors, indicating potential problems like transmission errors, hardware issues, or network congestion.
Policed discards	The frames discarded by the incoming packet match code.	The frames were discarded because they were not recognized or of interest. Usually, this field reports protocols that the Junos OS does not handle.
Drops	The number of packets dropped by the output queue of the I/O Manager application-specific integrated circuit (ASIC).	If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism.
L3 incompletes	The number of packets discarded due to the packets failing Layer 3 header checks.	This counter increments when the incoming packet fails Layer 3 (usually IPv4) checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment.
L2 channel errors	The errors that occur when the software could not find a valid logical interface (such as fe-1/2/3.0) for an incoming frame.	This error increments when, for example, a lookup for a virtual LAN (VLAN) fails.
L2 mismatch timeouts	The count of malformed or short packets.	The malformed or short packets cause the incoming packet handler to discard the frame and be unreadable.

Table 150: Major Gigabit Ethernet Counters (*Continued*)

Counter	Description	Reason for Increment
FIFO errors	The number of first in, first out (FIFO) errors in the receive direction as reported by the ASIC on the Physical Interface Card (PIC).	The value in this field should always be 0. If this value is not zero, cabling could be badly organized or the PIC could be broken.
Output Errors		
Errors	The sum of outgoing frame terminations and FCS errors.	The outgoing errors counter increments when frames sent from the device have improper terminations or FCS errors, usually caused by hardware faults, incorrect configurations, or transmission issues during the forwarding process.
Collisions	The number of Ethernet collisions.	The Gigabit Ethernet PIC supports only full-duplex operation, so this number should always remain 0. If it is incrementing, there is a software bug.
Drops	The 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.
Aged packets	The number of packets that remained in shared packet SDRAM for so long that the system automatically purged them.	The value in this field should never increment. If it increments, it is probably a software bug or broken hardware.
HS link FCS errors, FIFO errors	The number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.	The value in this field should always be 0. If it increments, either the FPC or the PIC is broken.

Miscellaneous Counters

Table 150: Major Gigabit Ethernet Counters (*Continued*)

Counter	Description	Reason for Increment
Input DA rejects	The number of packets that the filter rejected because the destination Media Access Control (MAC) address of the packet is not on the accept list.	It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad Address Resolution Protocol (ARP) entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting).
Output packet pad count	The number of packets that the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware.	Usually, padding is done only on small ARP packets, but some very small Internet Protocol (IP) packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist, or it is misconfigured.
Output packet error count	Number of packets with an indicated error that the filter was given to transmit.	These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.
CAM destination filters, CAM source filters	The number of entries in the content-addressable memory (CAM) dedicated to destination and source MAC address filters.	There can be up to 64 source entries. If source filtering is disabled, which is the default, the value for these fields should be 0.

SEE ALSO

[Understanding Interfaces on ACX Series Universal Metro Routers](#)

Troubleshoot: 10-Gigabit Ethernet Port Stuck in Down State

SUMMARY

Learn how to troubleshoot a 10-Gigabit Ethernet port stuck in a down state. Discover steps to diagnose and resolve issues with Dense Port Concentrator (DPC) or Packet Interface Card (PIC) states on Juniper Networks routers.

IN THIS SECTION

- [Problem | 473](#)
- [Cause | 473](#)
- [Solution | 473](#)

Problem

10-Gigabit Ethernet port is stuck in Dense Port Concentrator (DPC) or Packet Interface Card (PIC) down state.

Symptoms

The device has failed to initialize because the Ethernet port is down.

Cause

Disable and reenable the interface and then reset the transceiver and cable. If the interface remains in the down state, it may be stuck in DPC or PIC down state.

Does the router function normally after you disable and reenable the interface and then reset the transceiver and cable?

Yes:

The system is not stuck in DPC or PIC down state. Disable and reenable the interface or reset the transceiver to resolve the issue.

No:

The interface might be stuck in DPC or PIC down state.

Solution

To resolve the issue

From the above-mentioned diagnosis, you ascertain that the interface is stuck in DPC or PIC down state.

This is not a hardware related issue or problem. To resolve this issue, implement one of the following solutions:

- Reset the PIC.
- Toggle the framing mode.

1. In configuration mode, go to the [edit interfaces] hierarchy level.

```
user@host1# edit interfaces interface name
```

2. Toggle the framing mode. In the following configuration, WAN-PHY mode is toggled.

```
[edit interfaces interface-name is in the et-fpc/pic/port
user@host1# set framing wan-phy
user@host1# commit
user@host1# framing {
user@host1# wan-phy;
user@host1# }
user@host1# delete framing
user@host1# commit
```

3. Reset the PIC

```
user@host1# request chassis fpc slot x offline
user@host1# request chassis fpc slot x online
```

RELATED DOCUMENTATION

| [DPCs Supported on MX240, MX480, and MX960 Routers](#)

Verify Link and Transceivers with Pseudo Random Binary Sequence (PRBS) Test

SUMMARY

Learn how to verify links and transceivers using the Pseudo Random Binary Sequence (PRBS) test and perform bidirectional diagnostics with remote loopback. Discover essential steps for clearing interface statistics effectively.

IN THIS SECTION

- [How to Use Pseudo Random Binary Sequence \(PRBS\) Test for Link and Transceivers Verification | 475](#)
- [Configure Unidirectional Diagnostics | 477](#)
- [Configure Bidirectional Diagnostics with Remote Loopback | 479](#)
- [Clear the Interface Statistics | 483](#)

How to Use Pseudo Random Binary Sequence (PRBS) Test for Link and Transceivers Verification

The Pseudo Random Binary Sequence (PRBS) test enables you to verify link quality and transceiver operation. There are two diagnostic scenarios:

- **Bidirectional verification:** The test sequence starts from the local end, and the remote end is configured with loopback. Analyze the test pattern at both ends.
- **Unidirectional verification:** The test sequence starts from the local end, and the remote end analyzes the test pattern.

In the first case, bidirectional verification requires loopback support on the remote end.

Table 151: PRBS test on various MICs:

MIC Type	PRBS Test Engine	Description
JNP10003-LC2103	Packet Forwarding Engine (PFE)	The PFE for JNP10003-LC2103 is a specialized hardware module designed to handle high-performance packet forwarding, routing, and traffic management for the Juniper Networks JNP10003 system.

Table 151: PRBS test on various MICs: (Continued)

MIC Type	PRBS Test Engine	Description
JNP-MIC1	Packet Forwarding Engine (PFE)	The PFE for JNP-MIC1 is a hardware module designed to manage and accelerate packet forwarding, enabling high-speed data processing and routing for the Juniper Networks JNP Series Modular Interface Cards (MIC).
JNP-MIC1-MACSEC	External physical layer (PHY) device	External physical layer (PHY) device is not capable of passing the PRBS pattern originating from the packet forwarding engine. Hence, the PRBS test is enabled on external physical layer (PHY) device.

Table 152: PRBS test details supported on various interfaces:

Interface Type	Interface Name	Interface Lane Characteristics	Description
10-Gigabit ethernet interface	xe	The interface is supported with 1 lane of 10Gbps speed	The PRBS test is executed on each lane supported. Hence, the show interfaces prbs-stats displays data for one lane.
40-Gigabit ethernet interface	et	The interface is supported with 4 lanes of 10Gbps speed.	The PRBS test is executed on each lane supported. Hence, the show interfaces prbs-stats displays data for four lanes.
100-Gigabit ethernet interface	et	The interface is supported with 4 lanes of 25Gbps speed.	The PRBS test is executed on each lane supported. Hence, the show interfaces prbs-stats displays data for four lanes.

Configure Unidirectional Diagnostics

- Enable the transmission of pseudo-random binary sequence pattern on one end of the link (TX), and enable monitoring on the other end (RX).
- The PRBS statistics report on the receiving end reflects link quality.
- Collect and view the PRBS statistics.

1. Start a TX (direction 0) by issuing the following command:

```
[edit ]
user@host1> test interface et-0/1/2 prbs-test-start pattern-type 31 direction 0 flip 0
```

Check the link status by executing **show interfaces terse et-0/1/***:

Interface	Admin	Link	Proto	Local	Remote
et-0/1/2	up	down			

2. Start a RX (direction 1) by issuing the following command:

```
[edit ]
user@host2> test interface et-1/1/4 prbs-test-start pattern-type 31 direction 1 flip 0
```

Check the link status by executing **show interfaces terse et-1/1/4**:

Interface	Admin	Link	Proto	Local	Remote
et-1/1/4	up	down			

3. View the collected statistics at RX by issuing the following command:

```
[edit ]
user@host2> show interfaces interface-name prbs-stats
```

For example:

Checking PRBS statistics at RX:

```
user@host2> show interfaces et-1/1/4 prbs-stats
PRBS Statistics : Enabled

Lane 0 : State : Pass, Error count : 0
Lane 1 : State : Pass, Error count : 0
Lane 2 : State : Pass, Error count : 0
Lane 3 : State : Pass, Error count : 0
```

The PRBS test is successful, if the state is pass with error count 0.

4. Stop the PRBS statistics collection by issuing the following command:

```
user@host2> test interface interface-name prbs-test-stop direction 1
user@host1> test interface interface-name prbs-test-stop direction 0
```

For example:

To stop the PRBS at RX:

```
user@host2> test interface et-1/1/4 prbs-test-stop direction 1
```

Verify the statistics at TX by executing the following command:

```
show interfaces et-0/1/2 prbs-stats
PRBS Statistics : Disabled
```

To stop the PRBS at TX:

```
user@host1> test interface et-0/1/2 prbs-test-stop direction 0
```

After executing the command, you can check the link status at RX by executing **show interfaces terse et-1/1/4**:

Interface	Admin	Link	Proto	Local	Remote
et-1/1/4	up	up			

Check the link status at TX by executing **show interfaces terse et-0/1/2**:

Interface	Admin	Link	Proto	Local	Remote
et-0/1/2	up			up	

This command only disables the statistics collection and does not clear the statistics collected. To clear the collected statistics, issue the **clear interfaces statistics** command.

```
user@host1> clear interfaces statistics et-0/1/2
```

Configure Bidirectional Diagnostics with Remote Loopback

Configure loopback on the remote end of the link. Test pattern generation and analysis is performed on the local end.

The PRBS statistics report on the receiving end reflects link quality.

Following are the steps to collect and view the PRBS statistics:

1. Enable remote loopback.

```
user@host2> set interfaces et-1/1/4 gigether-options loopback-remote
```

You must disable Forward Error Correction (FEC), if loopback is configured in the router with JNP-MIC1 MIC at the remote end.

```
user@host2> show interfaces et-1/1/4 | display set
set interfaces et-1/1/4 gigether-options fec none
```

2. Start a TX (direction 0) by issuing the following command:

```
[edit]
user@host1> test interface et-0/1/2 prbs-test-start pattern-type 31 direction 0 flip 0
```

After executing the command, you can check the link status by executing `show interfaces terse et-0/1/*`:

Interface	Admin	Link	Proto	Local	Remote
et-0/1/2	up			down	

3. Start a RX (direction 1) by issuing the following command on the same host.

```
[edit]
user@host1> test interface et-0/1/2 prbs-test-start pattern-type 31 direction 1 flip 0
```

4. View the collected statistics at RX by issuing the following command:

```
[edit]
user@host1> show interfaces interface-name prbs-stats
```

For example:

Checking PRBS statistics at RX:

```
user@host1> show interfaces et-0/1/2 prbs-stats
PRBS Statistics : Enabled

Lane 0 : State : Pass, Error count : 0
Lane 1 : State : Pass, Error count : 0
Lane 2 : State : Pass, Error count : 0
Lane 3 : State : Pass, Error count : 0
```

The PRBS test is successful, if the state is pass with error count 0.

5. Stop the PRBS statistics collection by issuing the following command:

```
user@host1> test interface interface-name prbs-test-stop direction 1
user@host1> test interface interface-name prbs-test-stop direction 0
```

For example:

To stop the PRBS at RX:

```
user@host1> test interface et-0/1/2 prbs-test-stop direction 1
```

Verify the statistics at TX by executing the following command:

```
show interfaces et-0/1/2 prbs-stats
PRBS Statistics : Disabled
```

To stop the PRBS at TX:

```
user@host1> test interface et-0/1/2 prbs-test-stop direction 0
```

After executing the command, you can check the link status at RX by executing show interfaces terse et-0/1/2:

Interface	Admin	Link	Proto	Local	Remote
et-0/1/2	up			up	

This command only disables the statistics collection and does not clear the statistics collected. To clear the collected statistics, issue the clear interfaces statistics command.

```
user@host1> clear interfaces statistics et-0/1/2
```

Table 153: Interface Card Specific Differences

Behavior	JNP-MIC1-MACSEC MIC	JNP-MIC1 MIC
RX Latching to PRBS Signal	<p>If RX is not latching to any PRBS signal, then the "state" in the show interfaces <i>interface-name</i> prbs-stats displays as "Disabled" with Error count as 0.</p> <pre>user@host> test interface et-0/1/10 prbs-test-start pattern-type 31 direction 1 flip 0 user@host> show interfaces et-0/1/10 prbs-stats</pre> <p>PRBS Statistics : Enabled Lane 0 : State : Disabled, Error count : 0 Lane 1 : State : Disabled, Error count : 0 Lane 2 : State : Disabled, Error count : 0 Lane 3 : State : Disabled, Error count : 0</p>	<p>If RX is not latching to any PRBS signal, then the "state" in the show interfaces <i>interface-name</i> prbs-stats displays as failed with MAX error count.</p> <pre>user@host> test interface et-0/0/1 prbs-test-start pattern-type 31 direction 1 flip 0 user@host> show interfaces et-0/0/1 prbs-stats</pre> <p>PRBS Statistics : Enabled Lane 0 : State : Fail, Error count : 4294967295 Lane 1 : State : Fail, Error count : 4294967295 Lane 2 : State : Fail, Error count : 4294967295 Lane 3 : State : Fail, Error count : 4294967295</p>
Pattern Mismatch	Displays a "disabled" state for pattern mismatches and flips between TX and RX (flip supported).	Maximum errors due to pattern mismatches between TX and RX flips (flip not supported).
TX Interruptions	RX shows "Disabled" with Error Count 0; no need to restart RX when TX restarts.	RX shows a failed state with error counts; RX must be restarted if TX restarts.
Consecutive Starts	Behaves differently from JNP-MIC1 MIC when starting TX or RX consecutively without stopping earlier runs.	Shows distinct behavior when starting TX or RX consecutively without stopping earlier runs.

Table 153: Interface Card Specific Differences (Continued)

Behavior	JNP-MIC1-MACSEC MIC	JNP-MIC1 MIC
DFE Tuning	Does not require Decision Feedback Equalization (DFE) tuning.	Requires DFE tuning; will show errors if PRBS restarts without stopping earlier runs.
Forward Error Correction (FEC)	No specific mention of FEC settings.	FEC must be disabled if loopback is configured at the remote end. user@host> show interfaces et-1/1/1 display set set interfaces et-1/1/1 gigether-options loopback-remote set interfaces et-1/1/1 gigether-options fec none

Clear the Interface Statistics

The clear interface statistics command clears only the error counters and not the status, RX needs to be restarted to get the right status.

Following are the steps to clear the interface statistics:

1. Check the statistics at RX by issuing the following command:

```
[edit]
user@host2> show interfaces et-1/1/4 prbs-stats
PRBS Statistics : Enabled
  Lane 0 : State : Fail, Error count : 4294967295
  Lane 1 : State : Fail, Error count : 4294967295
  Lane 2 : State : Fail, Error count : 4294967295
  Lane 3 : State : Fail, Error count : 4294967295
```

Because only RX is started, there is no PRBS pattern and RX shows maximum error.

2. To clear the interface statistics, execute the following command:

```
[edit]
user@host2> clear interfaces statistics et-1/1/4
```

```
user@host2> show interfaces et-1/1/4 prbs-stats
PRBS Statistics : Enabled
  Lane 0 : State : Fail, Error count : 0
  Lane 1 : State : Fail, Error count : 0
  Lane 2 : State : Fail, Error count : 0
  Lane 3 : State : Fail, Error count : 0
```

Here the status still shows as fail, although the statistics displays the delta value. In this case, because both the current and previous values are INT_MAX, the delta value 0 is displayed.

Consider a scenario where the error count increments as the PRBS test is ongoing. In this case, the **show interfaces *interface-name* prbs-stats** shows the incremental error (delta value). Also after issuing **clear interfaces statistics et-1/1/4**, if the error count is updated, the **clear interfaces statistics** displays the incremental error too.

For Example:

```
[edit]
user@host2> show interfaces et-1/1/4 prbs-stats
PRBS Statistics : Enabled
  Lane 0 : State : Fail, Error count : 640
  Lane 1 : State : Fail, Error count : 647
  Lane 2 : State : Fail, Error count : 661
  Lane 3 : State : Fail, Error count : 596

[edit]
user@host2> show interfaces et-1/1/4 prbs-stats
PRBS Statistics : Enabled
  Lane 0 : State : Fail, Error count : 52
  Lane 1 : State : Fail, Error count : 65
  Lane 2 : State : Fail, Error count : 626
  Lane 3 : State : Fail, Error count : 132
```

As shown above, in the second instance the **show interfaces *interface-name* prbs-stats** command displays the "change" in error count. So, the total error count is $640 + 52 = 692$ for Lane 0.

SEE ALSO

[prbs-test-start](#)

[prbs-test-stop](#)

[show-interfaces-prbs-stats](#)

| clear-interfaces-statistics

9

PART

Configuration Statements and Operational Commands

- Junos CLI Reference Overview | **487**

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Read this guide to learn about the syntax and options that make up the statements and commands. Also understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- [Junos CLI Reference](#)

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- [Configuration Statements](#)
- [Operational Commands](#)