

Junos® OS

Ethernet Interfaces User Guide for Routing Devices





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 © 2024 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 | xiv

1 Ethernet Interfaces

Configuring Ethernet Interfaces | 2

Ethernet Interfaces Overview | 2

Initial Configuration of Ethernet Interfaces | 3

Configuring Ethernet Physical Interface Properties | 4

Configuring the Interface Speed on Ethernet Interfaces | 6

Configuring the Ingress Rate Limit | 8

Configuring the Link Characteristics on Ethernet Interfaces | 8

Configuring Multicast Statistics Collection on Ethernet Interfaces | 9

MAC Address Validation on Static Ethernet Interfaces Overview | 10

Configuring MAC Address Validation on Static Ethernet Interfaces | 11

Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router | 12

Example: Configuring Fast Ethernet Interfaces | 15

Example: Configuring Gigabit Ethernet Interfaces | 15

Flow Control for Ethernet Interfaces | 17

Understanding Flow Control | 17

Configuring Flow Control | 18

MAC Address Filtering and Accounting on Ethernet Interfaces | 19

Configuring MAC Address Filtering for Ethernet Interfaces | 20

Enabling Source Address Filtering | 20

Configuring MAC Address Filtering on PTX Series Packet Transport Routers | 22

Configuring MAC Address Accounting | 23

Management Ethernet Interfaces | 25

Management Ethernet Interface Overview | 25

Configuring a Consistent Management IP Address | 27

Configuring the MAC Address on the Management Ethernet Interface | 29

Power over Ethernet (PoE) on ACX Series Routers | 30

```
Understanding PoE on ACX Series Universal Metro Routers | 31
   Example: Configuring PoE on ACX2000 Routers | 35
       Requirements | 35
       Overview | 35
       Configuration | 36
       Verification | 39
   Example: Disabling a PoE Interface on ACX2000 Routers | 41
       Requirements | 41
       Overview | 42
       Configuration | 42
       Verification | 43
   Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers | 43
Point-to-Point Protocol over Ethernet (PPPoE) | 45
   PPPoE Overview | 45
   Configuring PPPoE | 50
       Overview | 50
       Setting the Appropriate Encapsulation on the PPPoE Interface | 51
       Configuring PPPoE Encapsulation on an Ethernet Interface | 52
       Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface | 52
       Configuring the PPPoE Underlying Interface | 53
       Identifying the Access Concentrator | 53
       Configuring the PPPoE Automatic Reconnect Wait Timer | 53
       Configuring the PPPoE Service Name | 54
       Configuring the PPPoE Server Mode | 54
       Configuring the PPPoE Client Mode | 54
       Configuring the PPPoE Source and Destination Addresses | 55
       Deriving the PPPoE Source Address from a Specified Interface | 55
       Configuring the PPPoE IP Address by Negotiation | 56
       Configuring the Protocol MTU PPPoE | 56
       Example: Configuring a PPPoE Server Interface on an M120 or M320 Router | 57
   Disabling the Sending of PPPoE Keepalive Messages | 57
   Verifying a PPPoE Configuration | 58
   Tracing PPPoE Operations | 59
       Configuring the PPPoE Trace Log Filename | 60
       Configuring the Number and Size of PPPoE Log Files | 60
```

```
Configuring Access to the PPPoE Log File | 61
       Configuring a Regular Expression for PPPoE Lines to Be Logged | 61
       Configuring the PPPoE Tracing Flags | 61
Configuring Aggregated Ethernet Interfaces | 63
Aggregated Ethernet Interfaces | 63
    What Are Aggregated Ethernet Interfaces? | 64
    Configuration Guidelines for Aggregated Ethernet Interfaces | 65
    Configure Aggregated Ethernet Interfaces | 67
    Mixed-Mode and Mixed-Rate Aggregated Ethernet Interfaces | 69
    Platform Support for Mixed Aggregated Ethernet Bundles | 69
    Configuration Guidelines for Mixed-Rate Aggregated Ethernet Links | 72
    Configure Mixed-Rate Aggregated Ethernet Interfaces | 73
    What Is Link Aggregation Control Protocol? | 74
    Configuration Guidelines for LACP | 75
    Configure LACP | 75
    Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 77
    Example: Configure Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces
       Across Aggregated Ethernet Member Links | 79
       Requirements | 83
       Overview | 83
       Verification | 84
    MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet
       Interfaces | 91
    What Is Enhanced LAG? | 92
Link Protection of Aggregated Ethernet Interfaces | 93
    Configuring Aggregated Ethernet Link Protection | 93
       Configuring Link Protection for Aggregated Ethernet Interfaces | 94
       Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces | 94
       Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link | 95
       Disabling Link Protection for Aggregated Ethernet Interfaces | 96
    Configuring Aggregated Ethernet Minimum Links | 96
    Example: Configuring Aggregated Ethernet Link Protection | 97
Scheduling on Aggregated Ethernet Interfaces | 98
```

Configuring Shared Scheduling on Aggregated Ethernet Interfaces | 98

```
Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection | 99
Load Balancing on Aggregated Ethernet Interfaces | 100
    Load Balancing and Ethernet Link Aggregation Overview | 101
    Understanding Aggregated Ethernet Load Balancing | 102
    Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data | 104
    Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | 107
    Configuring Adaptive Load Balancing | 108
    Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series
       Routers | 110
       Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview | 110
       Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers | 110
       Configuring Symmetrical Load Balancing on Trio-Based MPCs | 114
       Example Configurations | 116
    Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series
       Routers | 117
    Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on
       MX Series Routers | 120
    Example: Configuring Aggregated Ethernet Load Balancing | 123
       Example: Configuring Aggregated Ethernet Load Balancing | 123
Performance Monitoring on Aggregated Ethernet Interfaces | 141
    ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces Overview | 142
    Guidelines for Configuring Performance Monitoring Functionalities on Aggregated Ethernet
       Interfaces | 144
Periodic Packet Management | 145
    Understanding Periodic Packet Management on MX Series Routers | 146
    Configuring Periodic Packet Management on MX Series Routers | 147
       Enabling Centralized Periodic Packet Management | 147
Understanding Ethernet Link Aggregation on ACX Series Routers | 148
Gigabit Ethernet Interfaces
Configuring Gigabit Ethernet Interfaces | 158
Configuring 10-Gigabit Ethernet PICs | 158
    10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 159
    12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 163
    24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 166
```

```
P2-10G-40G-QSFPP PIC Overview | 167
   Configuring the P2-10G-40G-QSFPP PIC | 178
       Configuring the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 178
       Configuring the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 179
       Configuring the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 179
       Configuring the PIC at Port Group Level | 180
       Configuring Framing Mode on P2-10G-40G-QSFPP PIC | 181
   Example: Configuring the P2-10G-40G-QSFPP PIC | 182
       Requirements | 183
       Overview | 183
       Configuration | 183
   Framing Overview | 187
   Understanding WAN Framing | 188
   Configuring Ethernet Framing | 188
   Modes of Operation | 189
   Configuring Mixed-Rate Mode Operation | 190
   Configuring Line-Rate Mode on PICs Supporting Oversubscription | 192
   Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC | 192
   Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 194
   Gigabit Ethernet Notification of Link Down Alarm Overview | 197
   Notification of Link Down for Optics Options Overview | 198
   Configuring Gigabit Ethernet Notification of Link Down Alarm | 198
   Configuring Link Down Notification for Optics Options Alarm or Warning | 198
Configuring 40-Gigabit Ethernet PICs | 200
   40-Gigabit Ethernet PIC Overview | 200
   Configuring 40-Gigabit Ethernet PICs | 202
Configuring 100-Gigabit Ethernet MICs/PICs | 203
   100-Gigabit Ethernet Interfaces Overview | 204
   MPC3E MIC Overview | 207
   100-Gigabit Ethernet Type 4 PIC with CFP Overview | 208
   Configuring 100-Gigabit Ethernet Type 4 PIC With CFP | 211
   Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 215
   100-Gigabit Ethernet Type 5 PIC with CFP Overview | 218
   100-Gigabit Ethernet Interfaces Interoperability | 220
```

```
Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-
       CFP | 223
   Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and
       PD-1CE-CFP-FPC4 | 224
       Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-
           CFP | 224
   Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-
       PTX-2-100GE-CFP | 226
   Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP
       and PD-1CE-CFP-FPC4 | 227
       Configuring SA Multicast Bit Steering Mode on 100-Gigabit Ethernet PIC P1-
           PTX-2-100GE-CFP | 227
       Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC
           PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface | 229
Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | 231
   Transporting Legacy Traffic over Packet Switched Networks | 231
   Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH
       Traffic | 234
       Requirements for Configuration of the Smart-SFPs on MX Series Routers | 234
       Overview | 235
       Configuring the DS3 Smart SFP | 235
       Verification | 237
   Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy SDH
       Traffic | 240
       Requirements for Configuration of the Smart-SFPs on MX Series Routers | 241
       Overview | 241
       Configuring the STM1 Smart SFP | 241
       Verification | 243
Configuring Layer 2 Overhead Attribute in Interface Statistics | 247
   Accounting of the Layer 2 Overhead Attribute in Interface Statistics | 247
   Configuring Layer 2 Overhead Accounting in Interface Statistics | 250
       Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level | 251
   Verifying the Accounting of Layer 2 Overhead in Interface Statistics | 252
Configuring Gigabit Ethernet Policers | 255
   Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 256
   Configuring Gigabit Ethernet Policers | 258
```

```
Overview | 258
       Configuring a Policer | 259
       Specifying an Input Priority Map | 260
       Specifying an Output Priority Map | 261
       Applying a Policer | 262
       Configuring MAC Address Filtering | 263
       Example: Configuring Gigabit Ethernet Policers | 264
    Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 266
       Overview | 266
       Configuring a Policer | 268
       Applying a Policer | 269
       Example: Configuring and Applying a Policer | 269
Gigabit Ethernet Autonegotiation | 271
    Gigabit Ethernet Autonegotiation Overview | 271
    Configuring Gigabit Ethernet Autonegotiation | 272
       Configuring Gigabit Ethernet Autonegotiation with Remote Fault | 272
       Configuring Flow Control | 272
       Configuring Autonegotiation Speed on MX Series Routers | 272
       Displaying Autonegotiation Status | 273
Port Speed
Port Speed Overview | 281
Introduction to Port Speed | 281
Configure Port Speed | 288
Port Speed on ACX Routers | 294
Port Speed on ACX710 Router | 294
Port Speed on ACX5448 Router | 296
Port Speed on ACX7024 Router | 299
Port Speed on ACX7100 Router | 301
Port Speed on ACX7332 and ACX7348 Routers | 307
Port Speed on ACX7509 Router | 312
```

Port Speed on MX Routers | 321

Port Speed on MX204 | 321 Port speed on MX304 Router | 326 Port Speed on MPC7E-MRATE | 329 Port Speed on MPC7E-10G | 332 Port Speed on MIC-MRATE | 334 Port Speed on MX10003 MPC | 338 Port Speed on MIC-MACSEC-20GE | 344 Port Speed on MPC10E-10C-MRATE | 345 Port Speed on MPC10E-15C-MRATE | 350 Port Speed on MX2K-MPC11E | 356 Port Speed on MX10K-LC480 | 366 Port Speed on MX10K-LC2101 | 368 Port Speed on MX10K-LC9600 | 372 Port Speed on PTX Routers | 374 Port Speed on PTX10K-LC1201 | 375 Port Speed on PTX10001-36MR Router | 383 Port Speed on PTX10K-LC1202-36MR | 394 Port Speed on PTX10003 | 402 400G ZR and 400G ZR+ Optical Transceivers Overview | 410 400ZR and 400G OpenZR+ | 412 Licensing | 412 Application Selection | 413 Configure Application Selection | 415 Target Output Power | 418

Configure Target Output Power | 418

```
Enhanced Loopback Options | 420
Tunable Wavelength | 421
Signal Power Performance | 421
Configure 400ZR and 400G OpenZR+ Coherent Optics | 423
100GbE, 40GbE, and 10GbE Optical Transceivers
100GbE, 40GbE, and 10GbE | 426
Overview | 426
Supported Features | 428
Supported PICs | 432
Configure Optics | 433
Supported OTN and Optics Options | 438
   Supported OTN Options on ACX6360 and ACX5448 Routers | 438
   Supported OTN Options on MX Series Routers | 444
   Supported OTN Options on PTX Series Routers | 454
   Supported Optics Options on ACX6360 and ACX5448-D Routers | 465
   Supported Optics Options on PTX10008 and PTX10016 Series Routers | 469
Forward Error Correction and Bit Error Rate | 475
   Overview | 475
   Supported Forward Error Correction Modes | 480
   ODU Path Delay Measurement for Performance Monitoring | 481
       Overview | 482
       Enabling ODU Path Delay Measurement | 483
       Disabling ODU Path Delay Measurement | 485
| 486
| 487
| 487
| 487
| 487
```

Interface Mapping and Modulation format for PTX10K-LC1104 Line Card | 487

6

Troubleshooting Information

Monitoring and Troubleshooting Ethernet Interfaces | 492 Passive Monitoring on Ethernet Interfaces Overview | 492 Enabling Passive Monitoring on Ethernet Interfaces | 494 Link Degrade Monitoring Overview | 497 Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501 Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501 Monitor Fast Ethernet and Gigabit Ethernet Interfaces | 503 Display the Status of Fast Ethernet Interfaces | 503 Display the Status of Gigabit Ethernet Interfaces | 505 Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface | 507 Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface | 509 Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 513 Fiber-Optic Ethernet Interface Specifications | 515 Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 517 Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 517 Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface | 519 Create a Loopback | 520 Create a Physical Loopback for a Fiber-Optic Interface | 520 Create a Loopback Plug for an RJ-45 Ethernet Interface | 520 Configure a Local Loopback | 521 Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 523 Configure a Static Address Resolution Protocol Table Entry | 527 Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 532 Ping the Fast Ethernet or Gigabit Ethernet Interface | 533 Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 535 Diagnose a Suspected Circuit Problem | 538 Performing Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 538 Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 539 Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | 540

Create a Physical Loopback for a Fiber-Optic Interface | 541

Create a Loopback | 541

Configure a Local Loopback | 542 Verify That the Gigabit Ethernet Interface Is Up | 544 Configure a Static Address Resolution Protocol Table Entry | 549 Clear Gigabit Ethernet Interface Statistics | 551 Ping the Gigabit Ethernet Interface | 551 Check for Gigabit Ethernet Interface Error Statistics | 553 Diagnose a Suspected Circuit Problem | 557 Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 558 Configuring Loopback Testing | 558 Configuring BERT Testing | 560 Starting and Stopping a BERT Test | 565 Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters | 566 Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters | 566 Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm | 567 Fast Ethernet and Gigabit Ethernet Counters | 570 Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State | 573 Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | 575 Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | 576 Guidelines to perform Bidirectional Diagnostics using Remote Loopback | 580 Clearing the Interface Statistics | 583

Configuration Statements and Operational Commands

Unused | 587

Junos CLI Reference Overview | 589

About This Guide

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

RELATED DOCUMENTATION

Configuring Packet Optical networks with PTX Series Devices



Ethernet Interfaces

Configuring Ethernet Interfaces | 2

Configuring Aggregated Ethernet Interfaces | 63

Configuring Ethernet Interfaces

IN THIS CHAPTER

- Ethernet Interfaces Overview | 2
- Initial Configuration of Ethernet Interfaces | 3
- Flow Control for Ethernet Interfaces | 17
- MAC Address Filtering and Accounting on Ethernet Interfaces | 19
- Management Ethernet Interfaces | 25
- Power over Ethernet (PoE) on ACX Series Routers | 30
- Point-to-Point Protocol over Ethernet (PPPoE) | 45

Ethernet Interfaces Overview

Ethernet was developed in the early 1970s at the Xerox Palo Alto Research Center (PARC) as a data-link control layer protocol for interconnecting computers. It was first widely used at 10 megabits per second (Mbps) over coaxial cables and later over unshielded twisted pairs using 10Base-T. More recently, 100Base-TX (Fast Ethernet, 100 Mbps), Gigabit Ethernet (1 gigabit per second [Gbps]), 10-Gigabit Ethernet (10 Gbps), and 100-Gigabit Ethernet (100 Gbps) have become available.

Juniper Networks routers support the following types of Ethernet interfaces:

- Fast Ethernet
- Tri-Rate Ethernet copper
- 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)
- 100-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 Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet physical connections

RELATED DOCUMENTATION

Configuring MAC Address Filtering for Ethernet Interfaces | 20

Configuring Flow Control | 18

Configuring the Interface Speed on Ethernet Interfaces | 6

Initial Configuration of Ethernet Interfaces

IN THIS SECTION

- Configuring Ethernet Physical Interface Properties | 4
- Configuring the Interface Speed on Ethernet Interfaces | 6
- Configuring the Ingress Rate Limit | 8
- Configuring the Link Characteristics on Ethernet Interfaces | 8
- Configuring Multicast Statistics Collection on Ethernet Interfaces | 9
- MAC Address Validation on Static Ethernet Interfaces Overview | 10
- Configuring MAC Address Validation on Static Ethernet Interfaces | 11
- Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router | 12
- Example: Configuring Fast Ethernet Interfaces | 15
- Example: Configuring Gigabit Ethernet Interfaces | 15

Ethernet Interfaces are networking Interfaces that provide traffic connectivity. You can configure physical Interfaces as well as the logical Interfaces on your device. This topic discusses how to configure the physical properties of an Interface specific to Fast-Ethernet Interfaces, Gigabit-Ethernet Interfaces,

and aggregated Ethernet Interfaces. You can also use this topic for information on how to 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.

Configuring Ethernet Physical Interface Properties

To configure physical interface properties, for Fast Ethernet and Gigabit-Ethernet, DWDM interfaces, and other interfaces, complete the following steps:

1. To configure Fast Ethernet-specific physical interface properties, include the fastether-options statement at the [edit interfaces fe-fpc/pic/port] hierarchy level:

[edit interfaces fe-fpc/pic/port] user@host# set fastether-options;

NOTE: The speed statement applies to the management Ethernet interface (fxp0 or em0), the Fast Ethernet 12-port and 48-port Physical Interface Card (PIC) interfaces and the MX Series Tri-Rate Ethernet copper interfaces. The Fast Ethernet, fxp0, and em0 interfaces can be configured for 10 Mbps or 100 Mbps (10m | 100m). The MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps (10m | 100m | 1g). The 4-port and 8-port Fast Ethernet PICs support a speed of 100 Mbps only.

MX Series routers support Gigabit Ethernet automatic line sensing of MDI (Media Dependent Interface) and MDIX (Media Dependent Interface with Crossover) port connections. MDI is the Ethernet port connection typically used on network interface cards (NIC). MDIX is the standard Ethernet port wiring for hubs and switches. This feature allows MX Series routers to automatically detect MDI and MDIX connections and configure the router port accordingly. You can disable this feature by using the no-auto-mdix statement at the [edit interfaces ge-fpcl/picl/port] hierarchy level.

NOTE: Junos OS supports Ethernet host addresses with no subnets. This enables you to configure an Ethernet interface as a host address (that is, with a network mask of /32), without requiring a subnet. Such interfaces can serve as OSPF point-to-point interfaces, and MPLS is also supported.

2. To configure physical interface properties specific to Gigabit Ethernet and 10-Gigabit Ethernet, include the gigether-options statement at the [edit interfaces ge-fpc/pic/port] or [edit interfaces xe-fpc/pic/port] hierarchy level:

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

3. For 10-Gigabit Ethernet DWDM-specific physical interface properties, 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;
```

To configure Gigabit Ethernet IQ-specific physical interface properties, include the gigether-options statement at the [edit interfaces ge-fpc/pic/port] hierarchy level. These statements are supported on 10-Gigabit Ethernet IQ2 and IQ2-E PIC. Some of these statements are also supported on Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

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

4. To configure 10-Gigabit Ethernet physical interface properties, include the lan-phy or wan-phy statement at the [edit interfaces xe-fpc/pic/port framing] hierarchy level.

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

5. To configure OAM 802.3ah support for Ethernet interfaces, include the oam statement at the [edit protocols] hierarchy level.

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

6. To configure Gigabit Ethernet IQ-specific logical interface properties, include the input-vlan-map, output-vlan-map, layer2-policer, and vlan-tags statements at the [edit interfaces *interface-name* unit *logical-unit-*

number] hierarchy level or [edit logical-systems logical-system-name interfaces interface-name unit logicalunit-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;
```

7. To configure aggregated Ethernet-specific physical interface properties, include the aggregated-ether-options statement at the [edit interfaces aex] hierarchy level:

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

SEE ALSO

Example: Configuring Gigabit Ethernet Interfaces | 15

Configuring the Interface Speed on Ethernet Interfaces

For M Series and T Series Fast Ethernet 12-port and 48-port PIC interfaces, the management Ethernet interface (fxp0 or em0), and the MX Series Tri-Rate Ethernet copper interfaces, you can explicitly set the interface speed. The Fast Ethernet, fxp0, and em0 interfaces can be configured for 10 Mbps or 100 Mbps (10m | 100m). The MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps (10m | 100m | 1g). For information about management Ethernet interfaces and to determine the management Ethernet interface type for your router, see Understanding Management Ethernet Interfaces and Supported Routing Engines by Router. MX Series routers, with MX-DPC and Tri-Rate Copper SFPs, support 20x1 Copper to provide backwards compatibility with 100/10BASE-T and 1000BASE-T operation through an Serial Gigabit Media Independent Interface (SGMII) interface.

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

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

2. To configure the speed, include the speed statement at the [edit interfaces *interface-name*] hierarchy level.

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

NOTE:

- By default, the M Series and T Series routers management Ethernet interface autonegotiates whether to operate at 10 megabits per second (Mbps) or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the no-concatenate statement in the [edit chassis] configuration hierarchy.
- Starting with 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 10mmaximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.
- When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.
- If the link partner does not support autonegotiation, configure either Fast Ethernet port manually to match its link partner's speed and link mode. When the link mode is configured, autonegotiation is disabled.
- On MX Series routers with tri-rate copper SFP interfaces, if the port speed is negotiated to the configured value and the negotiated speed and interface speed do not match, the link will not be brought up.
- When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.
- Starting with Junos OS Release 11.4, half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the speed statement, you must include the link-mode full-duplex statement at the same hierarchy level.

SEE ALSO

Configuring the Ingress Rate Limit

On Fast Ethernet 8-port, 12-port, and 48-port PIC interfaces only, you can apply port-based rate limiting to the ingress traffic that arrives at the PIC.

To configure an ingress rate limit on a Fast Ethernet 8-port, 12-port, or 48-port PIC interface, include the ingress-rate-limit statement at the [edit interfaces *interface-name* fastether-options] hierarchy level:

```
[edit interfaces interface-name fastether-options]
ingress-rate-limit rate;
```

rate can range in value from 1 through 100 Mbps.

SEE ALSO

ingress-rate-limit

Configuring the Link Characteristics on Ethernet Interfaces

Full-duplex communication means that both ends of the communication can send and receive signals at the same time. *Half-duplex* is also bidirectional communication, but signals can flow in only one direction at a time.

By default, the router's management Ethernet interface, fxp0 or em0, autonegotiates whether to operate in full-duplex or half-duplex mode. Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet and 10-Gigabit Ethernet, the link partner must also be set to full duplex.

NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is fxp0. For T1600 and T4000 routers configured in a routing matrix, and TX Matrix Plus routers, the management Ethernet interface is em0.

NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the fxp0 management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the fxp0 management Ethernet interface so that the commands reference the em0 management Ethernet interface instead.

NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

NOTE: Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.

To explicitly 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]
link-mode (full-duplex | half-duplex);
```

NOTE: Starting in Junos OS release 16.1R7 and later, the link-mode configuration is not supported on 10-Gigabit Ethernet interfaces.

SEE ALSO

link-mode

Configuring Multicast Statistics Collection on Ethernet Interfaces

T Series and TX Matrix routers support multicast statistics collection on Ethernet interfaces in both ingress and egress directions. The multicast statistics functionality can be configured on a physical interface thus enabling multicast accounting for all the logical interfaces below the physical interface.

The multicast statistics information is displayed only when the interface is configured with the multicast-statistics statement, which is not enabled by default.

Multicast statistics collection requires at least one logical interface is configured with family inet and/or inet6; otherwise, the commit for multicast-statistics will fail.

The multicast in/out statistics can be obtained via interfaces statistics query through CLI and via MIB objects through SNMP query.

To configure multicast statistics:

Include the multicast-statistics statement at the [edit interfaces interface-name] hierarchy level.

An example of a multicast statistics configuration for a Ethernet interface follows:

```
[edit interfaces]
    ge-fpc/pic/port {
    multicast-statistics;
}
```

To display multicast statistics, use the show interfaces *interface-name* statistics detail command.

SEE ALSO

multicast-statistics

Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces

MAC Address Validation on Static Ethernet Interfaces Overview

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

MAC address validation is supported on AE, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces (with or without VLAN tagging) on MX Series routers only.

There are two types of MAC address validation that you can configure:

• Loose—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not support the MAC address of the tuple

Continues to forward packets when the source address of the incoming packet does not match any of the trusted IP addresses.

• Strict—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

SEE ALSO

Configure ARP Learning and Aging Options

Configuring MAC Address Validation on Static Ethernet Interfaces

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address. MAC address validation is supported on AE, Fast Ethernet, Gigabit Ethernet, and 10–Gigabit Ethernet interfaces (with or without VLAN tagging) on MX Series routers only.

To configure MAC address validation on static Ethernet Interfaces:

1. In the configuration mode, at the [edit] hierarchy level, configure the static Ethernet interface.

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

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

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

3. Configure MAC address validation on the static Ethernet Interface. You can specify the type of MAC address validation you require. Possible values are: Strict and Loose. You can also specify the interface address.

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

4. Configure the static ARP entry by specifying the IP address and the MAC address that are to be mapped. The IP address specified must be part of the subnet defined in the enclosing address statement. The MAC address must be specified as hexadecimal bytes in the following formats:

nnnn. nnnn or *nn: nn: nn: nn: nn: nn: nn* format. For instance, you can use either 0011.2233.4455 or 00:11:22:33:44:55.

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

SEE ALSO

family

mac-validate

Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router

The router internal Ethernet interface connects the Routing Engine with the router's packet forwarding components. The Junos OS automatically configures internal Ethernet interfaces. For TX Matrix Plus routers, the internal Ethernet interfaces are <code>ixgbe0</code> and <code>ixgbe1</code>. For T1600 routers configured in a routing matrix, the internal Ethernet interfaces are <code>bcm0</code> and <code>em1</code>. For more information about internal Ethernet interfaces, see Understanding Internal Ethernet Interfaces.

NOTE: Do not modify or remove the configuration for the internal Ethernet interface that the Junos OS automatically configures. If you do, the router will stop functioning.

The following example is a sequence of show interfaces commands issued in a Junos OS command-line interface (CLI) session with a TX Matrix Plus router in a routing matrix. In the example, the TX Matrix Plus router, which is also called the switch-fabric chassis (SFC), is known by the IP host name host-sfc-0 and contains redundant Routing Engines. The commands display information about the management Ethernet interface and both internal Ethernet interfaces configured on the Routing Engine to which you are currently logged in:

user@host-sfc-0> show interfaces em0 terse				
Interface	Admin Link Proto	Local	Remote	
em0	up up			
em0.0	up up inet	192.168.35.95/24		
user@host-sfc-0> show interfaces ixgbe0 terse				
Interface	Admin Link Proto	Local	Remote	

```
ixgbe0
                               up
                        up
ixgbe0.0
                                             10.34.0.4/8
                        up
                               up
                                    inet
                                             162.0.0.4/2
                                    inet6
                                             fe80::200:ff:fe22:4/64
                                             fec0::a:22:0:4/64
                                             0x22000004
                                    tnp
user@host-sfc-0> show interfaces ixgbe1 terse
Interface
                        Admin Link Proto
                                                                    Remote
                                             Local
ixgbe1
                               up
                        up
ixgbe1.0
                        up
                               up
                                    inet
                                             10.34.0.4/8
                                             162.0.0.4/2
                                    inet6
                                             fe80::200:1ff:fe22:4/64
                                             fec0::a:22:0:4/64
                                             0x22000004
                                    tnp
```

The following example is a sequence of show interfaces commands issued in a CLI session with a T1600 router in a routing matrix. In the example, the T1600 router, which is also called the line-card chassis (LCC), is known by the IP host name host-sfc-0-lcc-2 and contains redundant Routing Engines.

This T1600 router is connected to the routing matrix through a connection in the TXP-SIB-F13 in slot 2 of the SCC. The commands display information about the management Ethernet interface and both internal Ethernet interfaces configured on the Routing Engine to which you are currently logged in:

NOTE: In a routing matrix, the show interfaces command displays information about the current router only. If you are logged in to the TX Matrix Plus router, the show interfaces command output does not include information about any of the attached T1600 routers. To display interface information about a specific T1600 router in the routing matrix, you must first log in to that router.

The previous example shows a CLI session with the TX Matrix Plus router. To display interface information about the T1600 router known as host-sfc-0-lcc-2, first use the request routing-engine login command to log in to that LCC.

```
user@host-sfc-0> request routing-engine login lcc 2
--- JUNOS 9.6I built 2009-06-22 18:13:04 UTC
% cli
warning: This chassis is a Line Card Chassis (LCC) in a multichassis system.
warning: Use of interactive commands should be limited to debugging.
```

warning: Normal CLI access is provided by the Switch Fabric Chassis (SFC).

warning: Please logout and log into the SFC to use CLI.

user@host-sfc-0-lcc-2> show interfaces em0 terse

Interface Admin Link Proto Local Remote

em0 up up

em0.0 up up inet 192.168.35.117/24

user@host-sfc-0-lcc-2> show interfaces bcm0 terse

Interface Admin Link Proto Local Remote

bcm0 up up

bcm0.0 up up inet 10.1.0.5/8

129.0.0.5/2

inet6 fe80::201:ff:fe01:5/64

fec0::a:1:0:5/64

tnp 0x1000005

user@host-sfc-0-lcc-2> show interfaces em1 terse

Interface Admin Link Proto Local Remote

em1 up up

em1.0 up up inet 10.1.0.5/8

129.0.0.5/2

inet6 fe80::201:1ff:fe01:5/64

fec0::a:1:0:5/64

tnp 0x1000005

SEE ALSO

Understanding Internal Ethernet Interfaces

Example: Configuring Fast Ethernet Interfaces

The following configuration is sufficient to get a Fast Ethernet interface up and running. By default, IPv4 Fast Ethernet interfaces use Ethernet version 2 encapsulation.

SEE ALSO

Management Ethernet Interfaces | 25

Configuring Ethernet Interfaces

Example: Configuring Gigabit Ethernet Interfaces

The following configuration is sufficient to get a Gigabit Ethernet, Tri-Rate Ethernet copper, or 10-Gigabit Ethernet interface up and running. By default, IPv4 Gigabit Ethernet interfaces on MX Series, M Series, and T Series routers use 802.3 encapsulation.

```
[edit]
user@host# set interfaces ge-2/0/1 unit 0 family inet address local-address
user@host# show
interfaces {
    ge-2/0/1 {
        unit 0 {
            family inet {
                  address local-address;
            }
        }
    }
```

```
}
```

The M160, M320, M120, T320, and T640 2-port Gigabit Ethernet PIC supports two independent Gigabit Ethernet links.

Each of the two interfaces on the PIC is named:

```
ge-fpc/pic/[0.1]
```

Each of these interfaces has functionality identical to the Gigabit Ethernet interface supported on the single-port PIC.

SEE ALSO

Configuring Gigabit and 10-Gigabit Ethernet Interfaces

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501

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
14.2	Starting with 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 10mmaximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.
11.4	Starting with Junos OS Release 11.4, half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the speed statement, you must include the link-mode full-duplex statement at the same hierarchy level.

RELATED DOCUMENTATION

MAC Address Filtering and Accounting on Ethernet Interfaces | 19

Management Ethernet Interfaces | 25

Flow Control for Ethernet Interfaces

IN THIS SECTION

- Understanding Flow Control | 17
- Configuring Flow Control | 18

The MX, T, and PTX Series routers support IEEE 802.3X Ethernet PAUSE method of flow control. Flow control is enabled by default on all physical interfaces. This topic provides an overview of flow control for Ethernet Interfaces. It also describes how to explicitly enable flow control as well as disable flow control for Ethernet Interfaces.

Understanding Flow Control

IN THIS SECTION

IEEE 802.3X Ethernet PAUSE | 17

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of 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 it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

By default, all forms of flow control are enabled. You must explicitly enable flow control on interfaces to pause traffic.

IEEE 802.3X Ethernet PAUSE

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. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node
 A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses;
 then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes
 sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

Symmetric Flow Control

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

You configure symmetric flow control by including the flow-control statement at 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. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

SEE ALSO

flow-control

Configuring Flow Control

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

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

no-flow-control;

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

flow-control;

You can include these statements at the following hierarchy levels:

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

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

SEE ALSO

flow-control

Ethernet Interfaces Overview | 2

RELATED DOCUMENTATION

Understanding Priority-Based Flow Control

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

MAC Address Filtering and Accounting on Ethernet Interfaces

IN THIS SECTION

- Configuring MAC Address Filtering for Ethernet Interfaces | 20
- Configuring MAC Address Filtering on PTX Series Packet Transport Routers | 22

Configuring MAC Address Accounting | 23

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. This topic describes how to enable MAC address filtering and how to configure MAC address accounting.

Configuring MAC Address Filtering for Ethernet Interfaces

IN THIS SECTION

Enabling Source Address Filtering | 20

Enabling Source Address Filtering

On aggregated Ethernet interfaces, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), 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 following hierarchy levels:

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

NOTE: When you integrate a standalone T640 router into a routing matrix, the PIC media access control (MAC) addresses for the integrated T640 router are derived from a pool of MAC addresses maintained by the TX Matrix router. For each MAC address you specify in the configuration of a formerly standalone T640 router, you must specify the same MAC address in the configuration of the TX Matrix router.

Similarly, when you integrate a T1600 or T4000 router into a routing matrix, the PIC MAC addresses for the integrated T1600 or T4000 router are derived from a pool of MAC addresses maintained by the TX Matrix Plus router. For each MAC address you specify in the

configuration of a formerly standalone T1600 or T4000 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

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 following hierarchy levels:

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

You can specify the MAC address as nn: nn: nn: nn: nn: nn or nnnn . nnnnn, 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.

NOTE: The source-address-filter statement is not supported on Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router); instead, include the accept-source-mac statement. For more information, see "Configuring Gigabit Ethernet Policers" on page 258.

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

Source address filtering does not work when Link Aggregation Control Protocol (LACP) is enabled. This behavior is not applicable to T series routers and PTX Series Packet Transport Routers. For more information about LACP, see "Aggregated Ethernet Interfaces" on page 63.

NOTE: On untagged Gigabit Ethernet interfaces, you should not configure the source-address-filter statement at the [edit interfaces ge-fpclpiclport gigether-options] hierarchy level and the accept-source-mac statement at the [edit interfaces ge-fpclpiclport gigether-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 [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.

NOTE: The source-address-filter statement is not supported on MX Series routers with MPC4E (model numbers: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE); instead, include the accept-source-mac statement. For more information, see "Configuring Gigabit Ethernet Policers" on page 258.

RELATED DOCUMENTATION

```
source-address-filter

Aggregated Ethernet Interfaces | 63

Ethernet Interfaces Overview | 2
```

Configuring MAC Address Filtering on PTX Series Packet Transport Routers

This topic describes how to configure MAC filtering on PTX Series Packet Transport Routers. MAC filtering enables you to specify the MAC addresses from which the Ethernet interface can receive packets.

MAC filtering support on PTX Series Packet Transport Routers 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.

When you filter logical and physical interfaces, you can specify up to 1000 MAC source addresses per port.

To configure MAC source address filtering for a physical interface, include the source-filtering and source-address-filter statements at the [edit interfaces et-fpc/pic/port gigether-options] hierarchy level:

```
}
}
```

The source-address-filter statement configures which MAC source addresses are filtered. The specified physical interface drops all packets from the MAC source addresses you specify. You can specify the MAC address as *nn:nn:nn:nn:nn* where *n* is a decimal digit. To specify more than one address, include multiple *mac-address* options in the source-address-filter statement.

To configure MAC source address filtering for a logical interface, include the accept-source-mac statement at the [edit interfaces et-fpc/pic/port unit logical-unit-number] hierarchy level:

```
[edit interfaces]
et-x/y/z {
    gigether-options {
        source-filtering;
    }
    unit logical-unit-number {
        accept-source-mac {
            mac-address mac-address;
        }
    }
}
```

The accept-source-mac statement configures which MAC source addresses are accepted on the logical interface. You can specify the MAC address as *nn:nn:nn:nn:nn* where *n* is a decimal digit. To specify more than one address, include multiple mac-address *mac-address* options in the accept-source-mac statement.

After an interface filter is configured, there is an accounting entry that is associated with the MAC address filter. Counters accumulate if there are packets with matching MAC source addresses. You can use the show interfaces mac-database Junos OS CLI command to view the address count.

SEE ALSO

show interfaces mac-database

Configuring MAC Address Accounting

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), for Gigabit Ethernet DPCs on MX Series routers, for 100-Gigabit Ethernet Type 5 PIC with CFP, and for MPC3E, MPC4E, MPC5E, MPC5EQ, and MPC6E MPCs, you can configure whether source and destination MAC addresses are dynamically learned.

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:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile] mac-learn-enable;
```

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:

```
[edit interfaces aex aggregated-ether-options ethernet-switch-profile]
mac-learn-enable;
```

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.

NOTE: 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.

SEE ALSO

Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 256

Configuring Gigabit Ethernet Policers | 258

Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 266

Configuring a Policer Overhead

RELATED DOCUMENTATION

Aggregated Ethernet Interfaces | 63

Configuring Gigabit Ethernet Policers | 255

Management Ethernet Interfaces

IN THIS SECTION

- Management Ethernet Interface Overview | 25
- Configuring a Consistent Management IP Address | 27
- Configuring the MAC Address on the Management Ethernet Interface | 29

To connect to the router via the management port, use the management Ethernet interface. This topic provides you an overview of the management Ethernet Interface and describes how to configure the IP address and MAC address for the interface.

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.

Table 1 on page 26 summarizes the management interfaces typically used on Junos and Junos Evolved platforms. It's always a good idea to refer to the specific documentation for your platform to confirm details about its management interface.

Refer to the Product Documentation page for details on your platform.

Alternatively, refer to the Day One + quick start guide for your platform at: Day One + Guides.

Table 1: Typical Management Interfaces on Junos and Junos Evolved Platforms

Platform	Interface Name	Description
MX Series routers	fxp0	The fxp0 interface is typically an RJ-45 port on the Routing Engine.
EX Series switches	me0, vme	The me0 interface is typically an RJ-45 port on the Routing Engine.
		The vme interface is used when the device is part of a Virtual Chassis (VC), and
		is accessed via the me0 port. For consistency you can configure and use the
		vme interface on a stand alone switch.
QFX Series switches	em0, vme	The em0 interface is typically an RJ-45 port on the Routing Engine.
		The vme interface is used when the device is part of a Virtual Chassis (VC), and
		is accessed via the emo port. For consistency you can configure and use the
		vme interface on a stand alone switch.

Table 1: Typical Management Interfaces on Junos and Junos Evolved Platforms (Continued)

Platform	Interface Name	Description
SRX Security Gateways	fxp0, ge-0/0/0	The fxp0 interface is typically an RJ-45 port on the Routing Engine.
		On some SRX platforms the ge-0/0/0 interface is used as the management interface.
Platforms running Junos Evolved, for example, PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130, QFX5220, etc.	re0:mgmt-* and re1:mgmt-*	Junos Evolved platforms typically support two management interfaces per Routing Engine. The two Routing Engines are identified as re0 and re1. The RJ-45 copper port on each Routing Engine is indexed as 0, while the SFP fiber port is indexed as 1. For example, the RJ-45 management Ethernet port on Routing Engine 0 is typically named re0:mgmt-0.

SEE ALSO

Ethernet Interfaces Overview | 2

Initial Configuration of Ethernet Interfaces | 3

Configuring 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.

NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is fxp0. For TX Matrix Plus routers and T1600 or T4000 routers configured in a routing matrix, the management Ethernet interface is em0.

NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the fxp0 management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the fxp0 management Ethernet interface so that the commands reference the em0 management Ethernet interface instead.

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;
```

```
}
}
```

This feature is available on all routers that include dual Routing Engines. On the TX Matrix router, this feature is applicable to the switch-card chassis (SCC) only.

SEE ALSO

Ethernet Interfaces Overview | 2

Initial Configuration of Ethernet Interfaces | 3

Configuring the MAC Address on the Management Ethernet Interface

By default, the router's management Ethernet interface uses as its MAC address the MAC address that is burned into the Ethernet card.

NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is fxp0. For TX Matrix Plus routers and T1600 routers configured in a routing matrix, and TX Matrix Plus routers with 3D SIBs, T1600 routers, and T4000 routers configured in a routing matrix, the management Ethernet interface is em0.

NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the fxp0 management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the fxp0 management Ethernet interface so that the commands reference the em0 management Ethernet interface instead.

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.

To change the management Ethernet interface's MAC address, include the mac statement at the [edit interfaces fxp0] or [edit interfaces em0] hierarchy level:

```
[edit interfaces (fxp0 | em0)]
mac mac-address;
```

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

NOTE: If you integrate a standalone T640 router into a routing matrix, the PIC MAC addresses for the integrated T640 router are derived from a pool of MAC addresses maintained by the TX Matrix router. For each MAC address you specify in the configuration of a formerly standalone T640 router, you must specify the same MAC address in the configuration of the TX Matrix router.

Similarly, if you integrate a standalone T1600 router into a routing matrix, the PIC MAC addresses for the integrated T1600 router are derived from a pool of MAC addresses maintained by the TX Matrix Plus router. For each MAC address you specify in the configuration of a formerly standalone T1600 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

SEE ALSO

Ethernet Interfaces Overview | 2

Initial Configuration of Ethernet Interfaces | 3

mac

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 2

Initial Configuration of Ethernet Interfaces | 3

Power over Ethernet (PoE) on ACX Series Routers

IN THIS SECTION

- Understanding PoE on ACX Series Universal Metro Routers | 31
- Example: Configuring PoE on ACX2000 Routers | 35
- Example: Disabling a PoE Interface on ACX2000 Routers | 41
- Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers | 43

You can configure the ACX2000 Universal Metro Routers to enable the Power over Ethernet (PoE) ports to transfer both data and electrical power over a copper Ethernet LAN cable. This topic provides an overview of PoE support on ACX2000 routers and also provides information on how to configure, disable, and troubleshoot the PoE interface configured on the ACX2000 device.

Understanding PoE on ACX Series Universal Metro Routers

IN THIS SECTION

- ACX2000 PoE Specifications | 31
 - PoE Classes and Power Ratings | 32
- PoE Options | 33

Power over Ethernet (PoE) is the implementation of the IEEE 802.3af and IEEE 802.3at standards that allows both data and electrical power to pass over a copper Ethernet LAN cable.

Juniper Networks provides PoE on ACX2000 Universal Metro Routers that allows power delivery up to 65 W per PoE port. PoE ports transfer electrical power and data to remote devices over standard twisted-pair cables in an Ethernet network. Using the PoE ports, you can plug in devices that require both network connectivity and electrical power, such as voice over IP (VoIP) and wireless LAN access points.

You can configure the ACX2000 Universal Metro Router to act as a power sourcing equipment (PSE), supplying power to powered devices that are connected on designated ports.

This topic contains the following sections: :

ACX2000 PoE Specifications

Table 2 on page 32 lists the PoE specifications for the ACX2000 routers.

Table 2: PoE Specifications for the ACX2000 Routers

Specifications	For ACX2000 Universal Metro Routers	
Supported standards	 IEEE 802.3 AF IEEE 802.3 AT (PoE+) Legacy (pre-standards) 	
Supported ports	Supported on only two Gigabit Ethernet ports (ge-0/1/3 and ge-0/1/7).	
Total PoE power sourcing capacity	130 W	
Default per port power limit	32 W	
Maximum per port power limit	65 W	
Power management modes	 class—Power allocated for each interface can be configured. static—Power allocated for interfaces is based on the class of powered device connected. high-power—Power allocated for interfaces up to 65 W per port. 	

PoE Classes and Power Ratings

A powered device is classified based on the maximum power that it draws across all input voltages and operational modes. When class-based power management mode is configured on the ACX2000 routers, power is allocated taking into account the maximum power ratings defined for the different classes of devices.

Table 3 on page 33 lists the classes and their power ratings as specified by the IEEE standards.

Table 3: ACX2000 Universal Metro Router PoE Specifications

Class	Usage	Minimum Power Levels Output from PoE Port
0	Default	15.4 W
1	Optional	4.0 W
2	Optional	7.0 W
3	Optional	15.4 W
4	Reserved	Class 4 power devices are eligible to receive power up to 30 W according to the IEEE standards.

PoE Options

For ACX2000 Universal Metro Routers that support PoE ports, the factory default configuration enables PoE on the PoE-capable ports, with default settings in effect. You might not have to do any additional configuration if the default settings work for you. Table 4 on page 33 shows the PoE configuration options and their default settings for the PoE controller and for the PoE interfaces.

Table 4: PoE Configuration Options and Default Settings

Option	Default	Description
PoE Controller Option	ns	
guard-band	o w	Reserves up to 19 W power from the PoE power budget to be used in the case of a spike in PoE power consumption.

Table 4: PoE Configu	uration Options and Def	fault Settings (Continued)
Option	Default	Description
management	static	Sets the PoE power management mode for the router. The power management mode determines how power to a PoE interface is allocated: • class—Power allocated for each interface can be configured. • static—Power allocated for interfaces is based on the class of powered device connected. • high-power—Power allocated for interfaces up to 65 W per port.
Interface Options		
disable (Power over Ethernet)	Not included in default configuration	When included in the configuration, disables PoE on the interface. The interface maintains network connectivity but no longer supplies power to a connected powered device. Power is not allocated to the interface.
priority (Power over Ethernet)	low	Sets an interface's power priority to either low or high . If power is insufficient for all PoE interfaces, the PoE power to low-priority interfaces is shut down before power to high-priority interfaces is shut down. Among interfaces that have the same

assigned priority, the power priority is determined by port number, with lower-numbered ports having higher priority.

When included in the configuration, enables the logging of

power consumption records on an interface. Logging occurs every 5 minutes for 1 hour unless you specify a different value

for interval (Power over Ethernet) or duration.

SEE ALSO

telemetries

Power over Ethernet (PoE) User Guide for EX Series Switches

Not included in default

configuration

Example: Configuring PoE on ACX2000 Routers

IN THIS SECTION

- Requirements | 35
- Overview | 35
- Configuration | 36
- Verification | 39

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices. These ports allow you to plug in devices that need both network connectivity and electric power, such as voice over IP (VoIP) phones, wireless access points, and IP cameras.

This example shows how to configure PoE to deliver power up to 65 W on ACX2000 interfaces:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.2 or later for ACX Series routers
- An ACX2000 router that supports PoE

Before you configure PoE, be sure you have:

 Performed the initial router configuration. See ACX Series Autoinstallation Overview, Verifying Autoinstallation on ACX Series Universal Metro Routers, and Boot Sequence on Devices with Routing Engines for details.

Overview

This example consists of a router that has eight ports. Only two ports—ge-0/1/3 and ge-0/1/7—support PoE, which means they provide both network connectivity and electric power for powered devices such as VoIP telephones, wireless access points, and IP security cameras that require power up to 65 W. The remaining six ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers.

Table 5 on page 36 details the topology used in this configuration example.

Table 5: Components of the PoE Configuration

Property	Settings
Hardware	ACX2000 router with 8 Gigabit Ethernet ports: Two PoE interfaces (ge-0/1/3 and ge-0/1/7) and 6 non-PoE interfaces (ge-0/1/0, ge-0/1/1, ge-0/1/2, ge-0/1/4, ge-0/1/5, ge-0/1/6).
VLAN name	default
Connection to a wireless access point (requires PoE)	ge-0/1/7
Power port priority	high
Maximum power available to PoE port	65 W
PoE management mode	high-power
Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required)	ge-0/1/0 through ge-0/1/2
Unused ports (for future expansion)	ge-0/1/4 through ge-0/1/6

Configuration

IN THIS SECTION

Procedure | 37

Results | 38

To configure PoE on an ACX2000 router:

Procedure

CLI Quick Configuration

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

```
set poe management high-power guard-band 19
set poe interface ge-0/1/3 priority high maximum-power 65 telemetries
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure PoE:

1. Set the PoE management mode to high-power.

```
[edit]
user@host# set poe management high-power
```

NOTE:

- Set the PoE management mode to high-power only when the power requirement is more than 32 W and up to 65 W. If the power requirement is less than or equal to 32 W, then you do not need to set the PoE management mode to high-power.
- The default management mode is static. In this mode, the power sourcing equipment can deliver power up to 32 W.
- 2. Reserve power wattage in case of a spike in PoE consumption.

```
[edit]
user@host# set poe guard-band 19
```

3. Enable PoE.

```
[edit]
user@host# edit poe interface ge-0/1/3
```

4. Set the power port priority.

```
[edit poe interface ge-0/1/3]
user@host# set priority high
```

5. Set the maximum PoE power for a port.

```
[edit poe interface ge-0/1/3]
user@host# set maximum-power 65
```

NOTE: Set the maximum PoE power for a port only when the power requirement is more than 32 W and up to 65 W. If the power requirement is less than or equal to 32 W, then you do not need to configure the maximum PoE power.

6. Enable the logging of PoE power consumption.

```
[edit poe interface ge-0/1/3]
user@host# set telemetries
```

Results

In configuration mode, confirm your configuration by entering the show poe interface ge-0/1/3 command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
  user@host# show poe interface ge-0/1/3
  priority high;
  maximum-power 65;
telemetries;
```

If you are done configuring the device, enter commit in configuration mode.

Verification

IN THIS SECTION

- Verifying the Status of PoE Interfaces | 39
- Verifying the Telemetry Data (History) for the Specified Interface | 40
- Verifying PoE Global Parameters | 40

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

Verifying the Status of PoE Interfaces

Purpose

Verify that the PoE interfaces are enabled and set to the desired priority settings.

Action

In operational mode, enter the show poe interface ge-0/1/3 command.

```
user@host> show poe interface ge-0/1/3
```

PoE interface status:

PoE interface : ge-0/1/3
Administrative status : Enabled
Operational status : Powered-up

Power limit on the interface : 65 W Priority : High Power consumed : 6.6 W Class of power device : 0

Meaning

The show poe interface ge-0/1/3 command lists PoE interfaces configured on the ACX2000 router, with their status, priority, power consumption, and class.

Verifying the Telemetry Data (History) for the Specified Interface

Purpose

Verify the PoE interface's power consumption over a specified period.

Action

In operational mode, enter the show poe telemetries interface command.

For all records:

```
user@host> show poe telemetries interface ge-0/1/3 all

Interface Sl No Timestamp Power Voltage

1 Mon May 14 00:45:05 2012 14.2 W 53.9 V

2 Mon May 14 00:44:04 2012 14.2 W 53.9 V

3 Mon May 14 00:43:03 2012 14.2 W 53.9 V
```

For a specific number of records:

```
user@host> show poe telemetries interface ge-0/1/3 2

Interface Sl No Timestamp Power Voltage

1 Mon May 14 00:45:05 2012 14.2 W 53.9 V

2 Mon May 14 00:44:04 2012 14.2 W 53.9 V
```

Meaning

The telemetry status displays the power consumption history for the specified interface, provided telemetry has been configured for that interface.

Verifying PoE Global Parameters

Purpose

Verify global parameters such as guard band, power limit, and power consumption.

Action

In operational mode, enter the show poe controller command.

Controller	Maximum	Power	Guard	Management	Status	Lldp
index	power	consumption	band			Priority
0	130.0 W	14.2 W	0 W	high-power	UP	

Meaning

The show poe controller command lists the global parameters configured on the router.

SEE ALSO

Understanding PoE on ACX Series Universal Metro Routers | 31

Example: Disabling a PoE Interface on ACX2000 Routers

IN THIS SECTION

- Requirements | 41
- Overview | 42
- Configuration | 42
- Verification | 43

This example shows how to disable PoE on all interfaces or on a specific interface.

Requirements

Before you begin:

• Configure PoE on all interfaces. See "Example: Configuring PoE on ACX2000 Routers" on page 35.

Overview

In this example, you disable PoE on all interfaces and on a specific interface, which in this case is ge-0/1/3.

Configuration

IN THIS SECTION

Procedure | 42

Procedure

Step-by-Step Procedure

1. Disable PoE on all interfaces.

```
[edit]
user@host# set poe interface all disable
```

2. Commit the configuration.

Step-by-Step Procedure

1. Disable PoE on a specific interface.

```
[edit]
user@host# set poe interface ge-0/1/3 disable
```

2. Commit the configuration.

Verification

To verify the configuration is working properly, enter the show poe interface command.

user@host>	show poe int	erface				
Interface	Admin status	Oper Max status power	Priority	Power consumptio	Class	
ge-0/1/3	Disabled	Disabled 32.0W	Low	0.0W	0	
ge-0/1/3 ge-0/1/7	Disabled Disabled	Disabled 32.0W Disabled 32.0W	Low Low	0.0W 0.0W	0	

user@host> show poe interface ge-0/1/3

PoE interface status:

PoE interface : ge-0/1/3
Administrative status : Disabled
Operational status : Disabled
Power limit on the interface : 32.0 W
Priority : Low
Power consumed : 0.0 W
Class of power device : 0

SEE ALSO

Configuring PoE Interfaces on EX Series Switches

Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers

IN THIS SECTION

- Problem | 44
- Solution | 44

Problem

Description

A Power over Ethernet (PoE) interface is not supplying power to the powered device.

Solution

Check for the items shown in Table 6 on page 44.

Table 6: Troubleshooting a PoE Interface

Items to Check	Explanation
Is interface PoE enabled?	Only interfaces ge-0/1/3 and ge-0/1/7 can function as PoE ports.
Has PoE capability been disabled for that interface?	Use the show poe interface command to check PoE interface status.
Is the cable properly seated in the port socket?	Check the hardware.
Does the powered device require more power than is available on the interface?	Use the show poe interface command to check the maximum power provided by the interface.
If the telemetries option has been enabled for the interface, check the history of power consumption.	Use the show poe telemetries command to display the history of power consumption.

SEE ALSO

Monitoring and Troubleshooting PoE

Troubleshooting PoE Interfaces

RELATED DOCUMENTATION

Configuring PoE Interfaces on EX Series Switches

Monitoring and Troubleshooting PoE

Point-to-Point Protocol over Ethernet (PPPoE)

IN THIS SECTION

- PPPoE Overview | 45
- Configuring PPPoE | 50
- Disabling the Sending of PPPoE Keepalive Messages | 57
- Verifying a PPPoE Configuration | 58
- Tracing PPPoE Operations | 59

Use the Point-to-Point Protocol over Ethernet (PPPoE) encapsulation to connect multiple hosts on an Ethernet LAN to a remote site via 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 | 46
- PPPoE Stages | 47
- Optional CHAP Authentication | 49

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.

M120, M320, and MX Series routers can be configured as a PPPoE access concentrator server. To configure a PPPoE server on an M120, M320, or MX Series Ethernet *logical interface*, specify PPPoE encapsulation, include the pp0 statement for the pseudo PPPoE physical interface, and include the server statement in the PPPoE options under the logical interface.

NOTE: PPPoE encapsulation is not supported on M120, M320, or MX Series routers on an ATM2 IQ interface.

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.

This overview contains the following topics:

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:

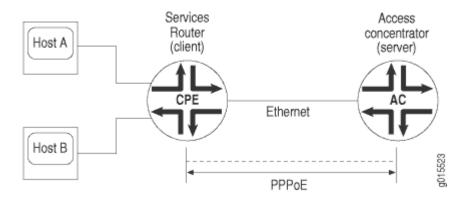
- If the interface is Fast Ethernet, use a PPPoE encapsulation.
- If the interface is ATM over ADSL, use a PPPoE over ATM encapsulation.

The PPPoE interface on M120 or M320 routers acting as a access concentrator can be a Gigabit Ethernet or 10-Gigabit Ethernet interface.

Ethernet Interface

The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 1 on page 47 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 1: 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.

NOTE: 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.

NOTE: 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 it 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.

For more information about CHAP, see *Configuring the PPP Challenge Handshake Authentication Protocol.*

SEE ALSO

Configuring the PPP Challenge Handshake Authentication Protocol

Evaluation Order for Matching Client Information in PPPoE Service Name Tables

Benefits of Configuring PPPoE Service Name Tables

Configuring PPPoE | 50

Disabling the Sending of PPPoE Keepalive Messages | 57

Configuring PPPoE Service Name Tables

Creating a Service Name Table

Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag

Configuring the Action Taken for the Any Service

Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag

Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information

Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name

Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client

Enabling Advertisement of Named Services in PADO Control Packets

Assigning a Service Name Table to a PPPoE Underlying Interface

Example: Configuring a PPPoE Service Name Table

Tracing PPPoE Operations | 59

Troubleshooting PPPoE Service Name Tables

Verifying a PPPoE Configuration | 58

Ethernet Interfaces User Guide for Routing Devices

Configuring PPPoE

IN THIS SECTION

- Overview | 50
- Setting the Appropriate Encapsulation on the PPPoE Interface | 51
- Configuring PPPoE Encapsulation on an Ethernet Interface | 52
- Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface | 52
- Configuring the PPPoE Underlying Interface | 53
- Identifying the Access Concentrator | 53
- Configuring the PPPoE Automatic Reconnect Wait Timer | 53
- Configuring the PPPoE Service Name | 54
- Configuring the PPPoE Server Mode | 54
- Configuring the PPPoE Client Mode | 54
- Configuring the PPPoE Source and Destination Addresses | 55
- Deriving the PPPoE Source Address from a Specified Interface | 55
- Configuring the PPPoE IP Address by Negotiation | 56
- Configuring the Protocol MTU PPPoE | 56
- Example: Configuring a PPPoE Server Interface on an M120 or M320 Router | 57

Overview

To configure PPPoE on an M120 or M320 Multiservice Edge Router or MX Series 5G Universal Routing Platform operating as an access concentrator, perform the following tasks:

- 1. Configure PPPoE encapsulation for an Ethernet interface.
- 2. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
- **3.** Optionally, configure the maximum transmission unit (MTU) of the interface.
- **4.** Configure the operational mode as server.
- 5. Configure the PPPoE interface address.
- **6.** Configure the destination PPPoE interface address.
- 7. Optionally, configure the MTU size for the protocol family.

- **8.** Starting in Junos OS Release 10.0, optionally, configure one or more PPPoE service name tables and the action taken for each service in the tables.
- **9.** Starting in Junos OS Release 12.3, optionally, disable the sending of PADS messages that contain certain error tags.

NOTE: Starting in Junos OS Release 10.4, when you configure a static PPPoE logical interface, you must include the pppoe-options subhierarchy at the [edit interfaces pp0 unit *logical-unit-number*] hierarchy level or at the [edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*] hierarchy level. If you omit the pppoe-options subhierarchy from the configuration, the commit operation fails.

Setting the Appropriate Encapsulation on the PPPoE Interface

For PPPoE on an Ethernet interface, you must configure encapsulation on the logical interface and use PPP over Ethernet encapsulation.

For PPPoE on an ATM-over-ADSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL logical interface, use PPPoE over AAL5 LLC encapsulation. LLC encapsulation allows a single ATM virtual connection to transport multiple protocols.

NOTE: PPPoE encapsulation is not supported on an M120 or M320 router on an ATM2 IQ interface.

When you configure a point-to-point encapsulation such as PPP on a physical interface, the physical interface can have only one logical interface (only one unit statement) associated with it.

To configure physical interface properties, include the encapsulation statement at the [edit interfaces interface-name] hierarchy level:

[edit interfaces *interface-name*] encapsulation ethernet-over-atm;

To configure logical interface encapsulation properties, include the encapsulation statement:

encapsulation ppp-over-ether;

You can include this statement at the following hierarchy levels:

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

Perform the task appropriate for the interface on which you are using PPPoE. For more information on how to configure PPoE encapsulation on an ethernet interface and on an ATM-over-ADSL interface, see "Configuring PPPoE Encapsulation on an Ethernet Interface" on page 52 and "Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface" on page 52.

Configuring PPPoE Encapsulation on an Ethernet Interface

Both the client and the server must be configured to support PPPoE. To configure PPPoE encapsulation on an Ethernet interface, include the encapsulation statement:

```
encapsulation ppp-over-ether;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]

Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface

To configure the PPPoE encapsulation on a ATM-over-ADSL interface, perform the following steps:

1. Include the encapsulation statement at the [edit interfaces *interface-name*] hierarchy level, and specify ethernet-over-atm:

```
[edit interfaces pp0]
encapsulation ethernet-over-atm;
```

2. Configure LLC encapsulation on the logical interface by including the encapsulation statement and specifying ppp-over-ether-over-atm-llc:

```
encapsulation ppp-over-ether-over-atm-llc;
```

You can include this statement at the following hierarchy levels:

• [edit interfaces pp0 unit *logical-unit-number*]

 $\bullet \quad [\text{edit logical-systems} \quad \textit{logical-system-name} \quad \text{interfaces pp0 unit} \quad \textit{logical-unit-number}] \\$

Configuring the PPPoE Underlying Interface

To configure the underlying Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface, include the underlying-interface statement:

```
underlying-interface interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number* pppoe-options]
- [edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]

Specify the logical Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface as the underlying interface—for example, at-0/0/1.0 (ATM VC), fe-1/0/1.0 (Fast Ethernet interface), or ge-2/0/0 (Gigabit Ethernet interface).

Identifying the Access Concentrator

When configuring a PPPoE client, identify the access concentrator by a unique name by including the access-concentrator statement:

```
access-concentrator name;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number* pppoe-options]
- [edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]

Configuring the PPPoE Automatic Reconnect Wait Timer

By default, after a PPPoE session is terminated, the session attempts to reconnect immediately. When configuring a PPPoE client, you can specify how many seconds to wait before attempting to reconnect, by including the auto-reconnect statement:

```
auto-reconnect seconds;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number* pppoe-options]
- [edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]

You can configure the reconnection attempt to occur in 0 through 4,294,967,295 seconds after the session terminates.

Configuring the PPPoE Service Name

When configuring a PPPoE client, identify the type of service provided by the access concentrator—such as the name of the Internet service provider (ISP), class, or quality of service—by including the service-name statement:

service-name *name*;

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number* pppoe-options]
- [edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options]

Configuring the PPPoE Server Mode

When configuring a PPPoE server, identify the mode by including the server statement:

server;

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number* pppoe-options]
- [edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]

Configuring the PPPoE Client Mode

When configuring a PPPoE client, identify the mode by including the client statement:

client;

You can include this statement at the following hierarchy levels:

• [edit interfaces pp0 unit *logical-unit-number* pppoe-options]

• [edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]

Configuring the PPPoE Source and Destination Addresses

When configuring a PPPoE client or server, assign source and destination addresses—for example, 192.168.1.1/32 and 192.168.1.2. To assign the source and destination address, include the address and destination statements:

```
address address {
    destination address;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces pp0.0 family inet]
- [edit logical-systems *logical-system-name* interfaces pp0.0 family inet]

Deriving the PPPoE Source Address from a Specified Interface

For a router supporting PPPoE, you can derive the source address from a specified interface—for example, the loopback interface, 100.0—and assign a destination address—for example, 192.168.1.2. The specified interface must include a logical unit number and have a configured IP address. To derive the source address and assign the destination address, include the unnumbered-address and destination statements:

```
unnumbered-address interface-name destination address;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces pp0.0 family inet]
- [edit logical-systems *logical-system-name* interfaces pp0.0 family inet]

Configuring the PPPoE IP Address by Negotiation

You can have the PPPoE client router obtain an IP address by negotiation with the remote end. This method might require the access concentrator to use a RADIUS authentication server. To obtain an IP address from the remote end by negotiation, include the negotiate-address statement:

```
negotiate-address;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0.0 family (inet | inet6 | mpls)]
- [edit logical-systems logical-system-name interfaces pp0.0 family (inet | inet6 | mpls)]

Configuring the Protocol MTU PPPoE

You can configure the maximum transmission unit (MTU) size for the protocol family. Specify a range from 0 through 5012 bytes. Ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. To set the MTU, include the mtu statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0.0 family (inet | inet6 | mpls)]
- [edit logical-systems logical-system-name interfaces pp0.0 family (inet | inet6 | mpls)]

You can modify the MTU size of the interface by including the mtu bytes statement at the [edit interfaces pp0] hierarchy level:

```
[edit interfaces pp0]
mtu bytes;
```

The default media MTU size used and the range of available sizes on a physical interface depends on the encapsulation used on that interface.

Example: Configuring a PPPoE Server Interface on an M120 or M320 Router

Configure a PPPoE server over a Gigabit Ethernet interface:

```
[edit interfaces]
ge-1/0/0 {
    vlan-tagging;
    unit 1 {
        encapsulation ppp-over-ether;
        vlan-id 10;
   }
}
pp0 {
    unit 0 {
        pppoe-options {
            underlying-interface ge-1/0/0.1;
            server;
        }
        ppp-options {
        family inet {
            address 22.2.2.1/32 {
                destination 22.2.2;
            }
        }
    }
}
```

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 2

Initial Configuration of Ethernet Interfaces | 3

Disabling the Sending of PPPoE Keepalive Messages

When configuring the client, you can disable the sending of keepalive messages on a logical interface by including the no-keepalives statement:

```
no-keepalives;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0 unit *logical-unit-number*]
- [edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number]

SEE ALSO

PPPoE Overview | 45

Ethernet Interfaces User Guide for Routing Devices

Verifying a PPPoE Configuration

IN THIS SECTION

- Purpose | **58**
- Action | 58

Purpose

You can use show commands to display and verify the PPPoE configuration.

Action

To verify a PPPoE configuration, you can issue the following operational mode commands:

- show interfaces at-fpc/pic/port extensive
- show interfaces pp0
- show pppoe interfaces
- show pppoe version
- show pppoe service-name-tables
- show pppoe sessions
- show pppoe statistics
- show pppoe underlying-interfaces

For more information about these operational mode commands, see CLI Explorer.

SEE ALSO

PPPoE Overview | 45

Ethernet Interfaces User Guide for Routing Devices

Tracing PPPoE Operations

IN THIS SECTION

- Configuring the PPPoE Trace Log Filename | 60
- Configuring the Number and Size of PPPoE Log Files | 60
- Configuring Access to the PPPoE Log File | 61
- Configuring a Regular Expression for PPPoE Lines to Be Logged | 61
- Configuring the PPPoE Tracing Flags | 61

The Junos OS trace feature tracks PPPoE operations and records events in a log file. The error descriptions captured in the log file provide detailed information to help you solve problems.

By default, nothing is traced. When you enable the tracing operation, the default tracing behavior is as follows:

- **1.** Important events are logged in a file called pppoed located in the /var/log directory. You cannot change the directory (/var/log) in which trace files are located.
- **2.** When the file pppoed reaches 128 kilobytes (KB), it is renamed pppoed.0, then pppoed.1, and finally pppoed.2, until there are three trace files. Then the oldest trace file (pppoed.2) is overwritten.

You can optionally specify the number of trace files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB). (For more information about how log files are created, see the System Log Explorer.)

By default, only the user who configures the tracing operation can access log files. You can optionally configure read-only access for all users.

To configure PPPoE tracing operations:

1. Specify that you want to configure tracing options.

[edit protocols pppoe]
user@host# edit traceoptions

- 2. (Optional) Configure the name for the file used for the trace output.
- 3. (Optional) Configure the number and size of the log files.
- 4. (Optional) Configure access to the log file.
- 5. (Optional) Configure a regular expression to filter logging events.
- **6.** (Optional) Configure flags to filter the operations to be logged.

Optional PPPoE traceoptions operations are described in the following sections:

Configuring the PPPoE Trace Log Filename

By default, the name of the file that records trace output for PPPoE is pppoed. You can specify a different name with the file option.

SEE ALSO

Tracing PPPoE Operations | 59

traceoptions (PPPoE)

Configuring the Number and Size of PPPoE Log Files

You can optionally specify the number of compressed, archived trace log files to be from 2 through 1000. You can also configure the maximum file size to be from 10 KB through 1 gigabyte (GB); the default size is 128 kilobytes (KB).

The archived files are differentiated by a suffix in the format .number.gz. The newest archived file is .0.gz and the oldest archived file is .(maximum number)-1.gz. When the current trace log file reaches the maximum size, it is compressed and renamed, and any existing archived files are renamed. This process repeats until the maximum number of archived files is reached, at which point the oldest file is overwritten.

For example, you can set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation, *filename*, reaches 2 MB, *filename* is compressed and renamed *filename*.0.gz, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename*.0.gz is renamed *filename*.1.gz and *filename* is compressed and renamed *filename*.0.gz. This process

repeats until there are 20 trace files. Then the oldest file, *filename*.19.gz, is simply overwritten when the next oldest file, *filename*.18.gz is compressed and renamed to *filename*.19.gz.

SEE ALSO

Tracing PPPoE Operations | 59

traceoptions (PPPoE)

Configuring Access to the PPPoE Log File

By default, only the user who configures the tracing operation can access the log files. You can enable all users to read the log file and you can explicitly set the default behavior of the log file.

SEE ALSO

Tracing PPPoE Operations | 59

traceoptions (PPPoE)

Configuring a Regular Expression for PPPoE Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including regular expressions to be matched.

SEE ALSO

Tracing PPPoE Operations | 59

traceoptions (PPPoE)

Configuring the PPPoE Tracing Flags

By default, no events are logged. You can specify which events and operations are logged by specifying one or more tracing flags.

To configure the flags for the events to be logged, configure the flags:

[edit protocols pppoe traceoptions] user@host# set flag authentication

SEE ALSO

Tracing PPPoE Operations | 59

RELATED DOCUMENTATION

PPPoE Overview | 45

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
12.3	Starting in Junos OS Release 12.3, optionally, disable the sending of PADS messages that contain certain error tags.
10.4	Starting in Junos OS Release 10.4, when you configure a static PPPoE logical interface, you must include the pppoe-options subhierarchy at the [edit interfaces pp0 unit <code>logical-unit-number</code>] hierarchy level or at the [edit logical-systems <code>logical-system-name</code> interfaces pp0 unit <code>logical-unit-number</code>] hierarchy level.
10.0	Starting in Junos OS Release 10.0, optionally, configure one or more PPPoE service name tables and the action taken for each service in the tables.

CHAPTER 2

Configuring Aggregated Ethernet Interfaces

IN THIS CHAPTER

- Aggregated Ethernet Interfaces | 63
- Link Protection of Aggregated Ethernet Interfaces | 93
- Scheduling on Aggregated Ethernet Interfaces | 98
- Load Balancing on Aggregated Ethernet Interfaces | 100
- Performance Monitoring on Aggregated Ethernet Interfaces | 141
- Periodic Packet Management | 145
- Understanding Ethernet Link Aggregation on ACX Series Routers | 148

Aggregated Ethernet Interfaces

SUMMARY

Learn about aggregated Ethernet interfaces (or Ethernet link aggregation), how to configure an aggregated Ethernet interface, LACP, and other supported features.

IN THIS SECTION

- What Are Aggregated Ethernet Interfaces? | 64
- Configuration Guidelines for Aggregated
 Ethernet Interfaces | 65
- Configure Aggregated EthernetInterfaces | 67
- Mixed-Mode and Mixed-Rate Aggregated
 Ethernet Interfaces | 69
- Platform Support for Mixed Aggregated
 Ethernet Bundles | 69
- Configuration Guidelines for Mixed-Rate
 Aggregated Ethernet Links | 72

- Configure Mixed-Rate Aggregated Ethernet
 Interfaces | 73
- What Is Link Aggregation Control Protocol? | 74
- Configuration Guidelines for LACP | 75
- Configure LACP | 75
- Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 77
- Example: Configure Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links | 79
- MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet Interfaces | 91
- What Is Enhanced LAG? | 92

What Are Aggregated Ethernet Interfaces?

IN THIS SECTION

Benefits | 64

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.

Configuration Guidelines for Aggregated Ethernet Interfaces

IN THIS SECTION

Platform Support for LAG | 65

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.
- You must not configure aggregated Ethernet for subscriber management by using the ether-options statement. If you do so, subscriber management does not work properly—there are issues with subscriber accounting and statistics. Use the gigether-options statement to configure aggregated Ethernet interfaces on the member link interfaces.
- You cannot configure simple filters on member link interfaces in an aggregated Ethernet bundle.
- You cannot configure any IQ-specific capabilities such as MAC accounting, VLAN rewrites, or VLAN
 queuing on member link interfaces in an aggregated Ethernet bundle.

Platform Support for LAG

Table 7 on page 66 lists the MX Series routers and the maximum number of interfaces per LAG and the maximum number of LAG groups they support. MX Series routers can support up to 64 interfaces per LAG.

Table 7: Maximum Interface Per LAG and Maximum LAGs per MX Router

MX Series Routers	Maximum Interfaces per LAG	Maximum LAG Groups
MX5, MX10, MX40, MX80, and MX104	16	Limited by the interface capacity. 80 on MX104.
MX150	10	10
MX240, MX480, MX204, MX304, MX960, MX10003, MX10008, MX10016, MX2010, and MX2020	64	128 (Before 14.2R1) 1000 (14.2R1 and later)

Table 8 on page 66 lists the PTX Series routers and the maximum number of interfaces per LAG and the maximum number of LAG groups they support. PTX Series routers can support up to 128 LAGs.

Table 8: Maximum Interface Per LAG and Maximum LAGs per PTX Router

PTX Series Routers	Maximum Interfaces per LAG	Maximum LAG Groups
PTX1000, PTX10002, and PTX10003, and PTX10008	64	128
PTX3000 and PTX5000	64	128
(Junos OS Evolved) PTX10008	64	1152

Table 9 on page 66 lists the ACX Series routers and the maximum number of interfaces per LAG and the maximum number of LAG groups they support.

Table 9: Maximum Interface Per LAG and Maximum LAGs per ACX Router

ACX Series Routers	Maximum Interfaces per LAG	Maximum LAG Groups
ACX7509	64	128

Configure Aggregated Ethernet Interfaces

Table 10 on page 67 describes the steps to configure aggregated Ethernet interfaces on your routing device.

Table 10: Aggregated Ethernet Interfaces Configuration

Configuration Step	Command
Step 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.	<pre>[edit chassis aggregated-devices ethernet] user@host# set device-count number</pre>
Step 2: Specify the members you want to include within the aggregated Ethernet bundle and add them individually. Aggregated interfaces are numbered from ae0 through ae4092.	[edit interfaces] user@host# set interface-name gigether-options 802.3ad aex
Step 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.	<pre>[edit interfaces] user@host# set aex aggregated-ether-options link-speed speed</pre>
Step 4: Specify the minimum number of links for the aggregated Ethernet interface (aex) —that is, the defined bundle— to be labeled <i>up</i> . By default, only one link must be up for the bundle to be labeled <i>up</i> . You cannot configure the minimum number of links and the minimum bandwidth at the same time. They are mutually exclusive.	[edit interfaces] user@host# set aex aggregated-ether-options minimum-links numbe
Step 5: (Optional) Specify the minimum bandwidth for the aggregated Ethernet links. You cannot configure link protection with minimum bandwidth. You cannot configure the minimum number of links and the minimum bandwidth at the same time. They are mutually exclusive.	[edit interfaces] user@host# set aex aggregated-ether-options minimum-bandwidth

Table 10: Aggregated Ethernet Interfaces Configuration (Continued)

Configuration Step	Command
Step 6: Specify an interface family and the IP address for the aggregated Ethernet bundle. Aggregated Ethernet interfaces can be VLAN-tagged or untagged. Packet tagging provides a logical way to differentiate traffic on ports which support multiple virtual local area network (VLAN). While you must configure aggregated Ethernet interfaces to receive tagged traffic, you must also configure aggregated Ethernet interfaces that can receive untagged traffic.	Tagged Interface [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
Step 7: (Optional) Configure your device to collect multicast statistics for the aggregated Ethernet interface. To view the multicast statistics, use the show interfaces statistics detail command. If you have not configured collection of multicast statistics, you cannot view the multicast statistics.	<pre>[edit interfaces] user@host# set aex multicast-statistics</pre>
Step 8: Verify and commit the configuration.	<pre>[edit interfaces] user@host# run show configuration user@host# commit</pre>
Step 9: (Optional) Delete an aggregated Ethernet Interface.	<pre>[edit] user@host# delete interfaces aex OR [edit] user@host# delete chassis aggregated-devices ethernet device- count</pre>

SEE ALSO

802.3ad

device-count

link-speed			
minimum-bandwi	dth		
minimum-links			
multicast-statistic	rs		

Mixed-Mode and Mixed-Rate Aggregated Ethernet Interfaces

IN THIS SECTION

Benefits | 69

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.

Platform Support for Mixed Aggregated Ethernet Bundles

Table 11 on page 69 lists the platforms and corresponding MPCs that support mixed-rate aggregated Ethernet bundles on MX Series routers.

For information about devices that support Mixed Aggregated Ethernet Bundles in Junos OS and Junos OS Evolved see Feature Explorer.

Table 11: Platform Support Matrix for Mixed-Rate Aggregated Ethernet Bundles on MX Series Routers

Supported MPCs	Supported Platform	Initial Release
16x10GE (MPC-3D-16XGE-SFPP)	MX240, MX480, MX960, MX2010, and MX2020	14.2R1

Table 11: Platform Support Matrix for Mixed-Rate Aggregated Ethernet Bundles on MX Series Routers *(Continued)*

Supported MPCs	Supported Platform	Initial Release
MPC1E (MX-MPC1-3D; MX-MPC1E-3D; MX-MPC-1-3D-Q; MX-MPC1E-3D-Q)	MX240, MX480, MX960, MX2010, and MX2020	14.2R1
MPC2E (MX-MPC2-3D; MX-MPC2E-3D; MX-MPC2-3D-Q; MX-MPC2-3D-Q; MX-MPC2-3D-EQ; MX-MPC2-3D-EQ; MX-MPC2-3D-P)	MX240, MX480, MX960, MX2010, and MX2020	14.2R1
MPC3E (MX-MPC3E-3D)	MX240, MX480, MX960, MX2010, and MX2020	14.2R1
MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE)	MX240, MX480, MX960, MX2010, and MX2020	14.2R1
MPC5E (6x40GE+24x10GE;6x40GE +24x10GEQ;2x100GE+4x10GE; 2x100GE+4x10GEQ)	MX240, MX480, MX960, MX2010, and MX2020	14.2R1
MPC6E (MX2K-MPC6E)	MX2010 and MX2020	14.2R1
MPC7E (Multi-Rate) (MPC7E-MRATE)	MX240, MX480, MX960, MX2010, and MX2020	15.1F4
MPC7E 10G (MPC7E-10G)	MX240, MX480, MX960, MX2010, and MX2020	15.1F5
MPC8E (MX2K-MPC8E)	MX2010 and MX2020	15.1F5
MPC9E (MX2K-MPC9E)	MX2010 and MX2020	15.1F5
MPC10E (MPC10E-15C-MRATE, MPC10E-10C-MRATE)	MX240, MX480, and MX960	19.1R1

Table 11: Platform Support Matrix for Mixed-Rate Aggregated Ethernet Bundles on MX Series Routers *(Continued)*

Supported MPCs	Supported Platform	Initial Release
MPC11E (MX2KE-MPC11E)	MX2010 and MX2020	19.3R2 and 20.1R1

Table 12 on page 71 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

Table 12: Platform Support Matrix for Mixed Aggregated Ethernet Bundles on T Series

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
10-Gigabit Ethernet LAN and WAN (WAN rate: OC192)	T640, T1600, T4000, and TX Matrix Plus routers	• T4000 FPC5 (T4000- FPC5-3D)	 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (PF-24XGE-SFPP) 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)
		• Enhanced Scaling FPC3 (T640-FPC3-ES)	10-Gigabit Ethernet PIC with XENPAK (PC-1XGE-XENPAK)
		 Enhanced Scaling FPC4 (T640-FPC4-ES) Enhanced Scaling FPC4-1P (T640- FPC4-1P-ES) T1600 Enhanced Scaling FPC4 (T1600- FPC4-ES) 	 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP) 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)
40-Gigabit Ethernet, 100- Gigabit Ethernet	T4000 and TX Matrix Plus routers	• T4000 FPC5 (T4000- FPC5-3D)	100-Gigabit Ethernet PIC with CFP (PF-1CGE-CFP)

Table 12: Platform Support Matrix for Mixed Aggregated Ethernet Bundles on T Series (Continued)

Rate and Mode	Supported Platform	Supported FPCs	Supported PICs
	T640, T1600, T4000, and TX Matrix Plus routers	 Enhanced Scaling FPC4 (T640-FPC4-ES) Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES) T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES) 	 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4) NOTE: This PIC is available packaged only in an assembly with the T1600-FPC4-ES FPC. 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP)

Configuration Guidelines for 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 MX Series, 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.
- Packets are dropped when the total throughput of the hash flow exiting a member link (or the
 throughput of multiple hash flows exiting a single member link) exceeds the link speed of the
 member link. This can happen when the egress member link changes because of a link failure and the
 hash flow switches to a member link of speed that is less than the total throughput of the hash flow.
- 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

Table 13 on page 73 describes the steps to configure mixed-rate aggregated Ethernet bundle on your device.

Table 13: Mixed-Rate Aggregated Ethernet Configuration

Configuration Step	Command
Step 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.	<pre>[edit chassis aggregated-devices ethernet] user@host# set device-count number</pre>
Step 2: Specify the members you want to include within the aggregated Ethernet bundle. Aggregated interfaces are numbered from ae0 through ae4092.	<pre>[edit interfaces] user@host# set interface-name gigether-options 802.3ad aex</pre>
Step 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.	<pre>[edit interfaces] user@host# set aex aggregated-ether-options link-speed mixed</pre>
You cannot configure the minimum number of links for the aggregated Ethernet bundle to be labeled <i>up</i> , when you configure the link speed as mixed.	
Step 4: Specify the minimum bandwidth for the aggregated Ethernet links. You cannot configure link protection with the minimum bandwidth.	<pre>[edit interfaces] user@host# set aex aggregated-ether-options minimum- bandwidth</pre>

Table 13: Mixed-Rate Aggregated Ethernet Configuration (Continued)

Configuration Step	Command
Step 5: Verify and commit the configuration.	<pre>[edit interfaces] user@host# run show configuration user@host# commit</pre>

SEE ALSO

802.3ad

device-count

link-speed

minimum-bandwidth

What Is Link Aggregation Control Protocol?

IN THIS SECTION

Benefits | 75

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 an LACP exchange. The partner is the remote interface in an 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-There is 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.

Configuration Guidelines for 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. If you use link aggregation to link together the ports of a PIC that supports up to 16 forwarding classes with a PIC that supports up to 8 forwarding classes, 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), the commit operation throws an error.
- When you enable a device to process packets received on a member link irrespective of the LACP state if the state of the aggregated Ethernet bundle is up (by using the accept-data statement), then the device does not process the packets as 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.

NOTE: On EX2300 and EX3400 switches, the LACP protocol must be configured with a periodic SLOW timer to prevent flaps during CPU intensive operations events such as routing engine switchover, interface flaps, and exhaustive data collection from the packet forwarding engine.

Configure LACP

Table 14 on page 76 describes the steps to configure LACP on an aggregated Ethernet interface.

Table 14: LACP Configuration

Configuration Step	Command
Step 1: Specify the LACP transmission mode - active or passive.	<pre>[edit interfaces interface-name aggregated-ether-options] user@host# set lacp active user@host# set lacp passive</pre>
Step 2: Specify the interval at which the interfaces send LACP packets. When you configure different intervals for the active and passive interfaces, the <i>actor</i> transmits the packets at the rate configured on the <i>partner's</i> interface.	<pre>[edit interfaces interface-name aggregated-ether-options lacp] user@host# set periodic interval</pre>
Step 3: Configure the LACP system identifier. 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.	<pre>[edit interfaces interface-name aggregated-ether-options lacp] user@host# set system-id system-id</pre>
Step 4: Configure the LACP system priority at the Aggregated Ethernet interface level. This system priority takes precedence over the priority value 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.	<pre>[edit interfaces interface-name aggregated-ether-options lacp] user@host# set system-priority system-priority</pre>
Step 5: (Optional) Configure the LACP administrative key. You must configure MC-LAG to configure this option. For more information on MC-LAG, see Understanding Multichassis Link Aggregation Groups.	<pre>[edit interfaces interface-name aggregated-ether-options lacp] user@host# set admin-key number</pre>

Table 14: LACP Configuration (Continued)

Configuration Step	Command
Step 6: 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.	<pre>[edit interfaces interface-name aggregated-ether-options lacp] user@host# set hold-time timer-value</pre>
Step 7: Configure the device to process packets received on a member link irrespective of the LACP state if the aggregated interface status is up.	<pre>[edit interfaces interface-name aggregated-ether-options lacp] user@host# set accept-data</pre>
Step 8: Verify and commit the configuration.	<pre>[edit interfaces interface-name aggregated-ether-options lacp] user@host# run show configuration user@host# commit</pre>

SEE ALSO

hold-time		
periodic		
system-id		
system-priority		

Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links

IN THIS SECTION

Benefits | 78

By default, 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 based on the hashing algorithm. Egress policy is distributed between individual member interface schedulers or policers 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.

Example: Configure Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links

IN THIS SECTION

- Requirements | 83
- Overview | 83
- Verification | 84

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 ge-0/0/3 apply-groups-except INTF gigether-options 802.3ad ae10
distribution-list dl1
user@host# set interfaces ge-0/0/4 apply-groups-except INTF gigether-options 802.3ad ae10
distribution-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. Assign each member of the aggregated Ethernet bundle to a different distribution list.

```
[edit]
user@host# set interfaces ge-0/0/3 apply-groups-except INTF gigether-options 802.3ad ae10
distribution-list dl1
[edit]
user@host# set interfaces ge-0/0/4 apply-groups-except INTF gigether-options 802.3ad ae10
distribution-list dl2
```

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

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

4. 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 ge-0/0/3
apply-groups-except INTF;
gigether-options {
    802.3ad {
        ae10;
        distribution-list dl1;
    }
}
```

```
user@host# show interfaces ge-0/0/4
apply-groups-except INTF;
gigether-options {
    802.3ad {
        ae10;
        distribution-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:

- Junos OS Release 16.1 and later 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 logical interface unit numbers of the aggregated Ethernet bundle ending in an odd number are assigned to the

distribution list dl1 as the primary list, and those ending in an even number are assigned the distribution list dl2 as the primary list.

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 | 84

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 d/1 (ge-0/0/3) and those ending in an even number being assigned to the distribution list d/2 (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 :
   Input packets:
                                     0
   Output packets:
                                     0
 Dropped traffic statistics due to STP State:
  Input bytes :
  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
                                                         0
```

2	_			_		_
gress queues: 8 supported, 4 in use queue counters: Queued packets Transmitted packets Dropped packets 0	2		0	0		0
New Counters Queued packets Transmitted packets Dropped packets			-	0		0
0				stad as I t	D I	-T. 4
1	-	Queued packe			Dropped pa	
2						
Mapped forwarding classes	•					
### Apped forwarding classes 0						•
0 best-effort 1 expedited-forwarding 2 assured-forwarding 3 network-control ogical 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 0 0 Output: 2 0 92 0 Adaptive Statistics: Adaptive Adjusts: 0 Adaptive Updates: 0 Link: ge-0/0/3.101 Input: 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input: 0 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 0 ge-0/0/4.101 0 0 0 0 0 List-Type Status Primary Active		и				0
1 expedited-forwarding 2 assured-forwarding 3 network-control ogical 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 0 Output: 2 0 92 0 Adaptive Statistics: Adaptive Adjusts: 0 Adaptive Updates: 0 Link: ge-0/0/3.101 Input: 0 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input: 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 0 0 0 0 0 Output: 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	-		rding class	es		
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 0 0 Output: 2 0 92 0 Adaptive Statistics: Adaptive Scans: 0 Adaptive Updates: 0 Link: ge-0/0/3.101 Input: 0 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input: 0 0 0 0 0 0 Output: 0 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 0 ge-0/0/4.101 0 0 0 0 0 List-Type Status Primary Active						
Ogical 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:						
Ogical 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						
Description: matched odd Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.101] Encapsulation: ENET2 Statistics	3	network-cont	101			
Bundle: Input: 0 0 0 0 0 Output: 2 0 92 0 Adaptive Statistics: Adaptive Adjusts: 0 Adaptive Updates: 0 Link: ge-0/0/3.101 Input: 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input: 0 0 0 0 0 Output: 0 0 0 0 Output: 0 0 0 0 0 List-Type Status Primary Active	Description: ma	atched odd				
Input: 0 0 0 0 0 Output: 2 0 92 0 Adaptive Statistics: Adaptive Adjusts: 0	Statistics	Packets	pps	Bytes	bps	
Output: 2 0 92 0 Adaptive Statistics: 8 4daptive Adjusts: 0 <td< td=""><td>Bundle:</td><td></td><td></td><td></td><td></td><td></td></td<>	Bundle:					
Adaptive Statistics: Adaptive Adjusts: 0 Adaptive Scans : 0 Adaptive Updates: 0 Link: ge-0/0/3.101 Input: 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input: 0 0 0 0 0 Output: 0 0 0 0 Output: 0 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 0 ge-0/0/4.101 0 0 0 0 0 0 List-Type Status Primary Active	Input :	0	0	0	0	
Adaptive Adjusts: 0 Adaptive Scans : 0 Adaptive Updates: 0 Link: ge-0/0/3.101 Input : 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input : 0 0 0 0 0 Output: 0 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 0 ge-0/0/4.101 0 0 0 0 0 List-Type Status Primary Active	Output:	2	0	92	0	
Adaptive Scans : 0 Adaptive Updates: 0 Link: ge-0/0/3.101 Input : 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input : 0 0 0 0 0 Output: 0 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 0 ge-0/0/4.101 0 0 0 0 0 List-Type Status Primary Active	Adaptive Statis	stics:				
Adaptive Updates: 0 Link: ge-0/0/3.101 Input: 0 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input: 0 0 0 0 0 0 Output: 0 0 0 0 0 Output: 0 0 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 0 ge-0/0/4.101 0 0 0 0 0 List-Type Status Primary Active	Adaptive Ad	djusts: 0				
Link: ge-0/0/3.101 Input: 0 0 0 0 0 Output: 2 0 92 0 ge-0/0/4.101 Input: 0 0 0 0 0 Output: 0 0 0 0 0 Output: 0 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 0 ge-0/0/4.101 0 0 0 0 0 List-Type Status Primary Active	Adaptive So	cans : 0				
Input :	Adaptive Up	odates: 0				
Input: 0 0 9 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	Link:					
Output: 2 0 92 0 ge-0/0/4.101 Input: 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	ge-0/0/3.101					
ge-0/0/4.101 Input: 0 0 0 0 Output: 0 0 0 0 Aggregate member links: 2 Marker Statistics: Marker Rx ge-0/0/3.101 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	Input :	0	0	0	0	
Input: 0 0 0 0 0 0 Output: 0 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 0 ge-0/0/4.101 0 0 0 0 0 List-Type Status Primary Active	•	2	0	92	0	
Output: 0 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 0 ge-0/0/4.101 0 0 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 0 ge-0/0/4.101 0 0 0 0 0 0 List-Type Status Primary Active	· ·	0	0	0	0	
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	Output:	0	0	0	0	
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						
ge-0/0/3.101 0 0 0 0 0 0 ge-0/0/4.101 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	Aggregate membe	er links: 2				
ge-0/0/4.101 0 0 0 0 List-Type Status Primary Active	Marker Statisti	cs: Marker Rx	Resp Tx	Unknown Rx	Illegal Rx	
List-Type Status Primary Active	ge-0/0/3.101	0	0	0	0	
Primary Active	ge-0/0/4.101	0	0	0	0	
	List-Type	Status				
Interfaces:	Primary	Active				
	Int	cerfaces:				

```
ge-0/0/4
                              Up
  List-Type
                  Status
  Backup
                  Waiting
              Interfaces:
              ge-0/0/3
                              Up
                  Status
  List-Type
  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
                                                               bps
                                   pps
                                               Bytes
  Bundle:
                          0
                                                                 0
      Input:
                                                    0
      Output:
                                                 296
                                                                 0
  Adaptive Statistics:
      Adaptive Adjusts:
     Adaptive Scans :
     Adaptive Updates:
  Link:
    ge-0/0/3.102
     Input:
                                                   0
                                                                 0
     Output:
                                                 296
    ge-0/0/4.102
     Input:
                          0
                                     0
                                                   0
                                                                 0
      Output:
                                     0
                                                            Illegal Rx
  Marker Statistics:
                       Marker Rx
                                     Resp Tx
                                               Unknown Rx
    ge-0/0/3.102
                               0
                                           0
                                                         0
                                                                      0
                                           0
    ge-0/0/4.102
                                                         0
                                                                      0
  List-Type
                  Status
  Primary
                  Active
              Interfaces:
              ge-0/0/3
                              Up
                  Status
  List-Type
  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
                          3
                                     0
                                                 194
                                                                 0
      Output:
  Adaptive Statistics:
      Adaptive Adjusts:
     Adaptive Scans :
                                 0
     Adaptive Updates:
  Link:
    ge-0/0/3.103
     Input:
                          0
                                     0
                                                   0
                                                                 0
     Output:
                                                 194
    ge-0/0/4.103
     Input :
                          0
                                     0
                                                   0
                                                                 0
                                                    0
      Output:
  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
  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:
                          0
                                                   0
                                                                 0
      Input:
      Output:
                          2
                                     0
                                                  92
                                                                 0
  Adaptive Statistics:
      Adaptive Adjusts:
                                 0
      Adaptive Scans :
                                 0
     Adaptive Updates:
  Link:
    ge-0/0/3.104
                                                   0
                                                                 0
     Input:
                          0
                                     0
                          2
      Output:
                                                   92
                                                                 0
    ge-0/0/4.104
                          0
                                     0
                                                    0
                                                                 0
     Input:
                                                    0
      Output:
  Marker Statistics:
                       Marker Rx
                                     Resp Tx
                                               Unknown Rx
                                                            Illegal Rx
    ge-0/0/3.104
                               0
                                           0
                                                         0
                                                                      0
                                           0
                                                         0
                                                                      0
    ge-0/0/4.104
  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
                                                              bps
                                   pps
                                               Bytes
  Bundle:
                                                                0
     Input:
                                                   0
     Output:
                          0
                                     0
                                                   0
                                                                0
  Adaptive Statistics:
      Adaptive Adjusts:
     Adaptive Scans :
     Adaptive Updates:
  Link:
    ge-0/0/3.32767
     Input :
                                                                0
                          0
                                     0
                                                   0
      Output:
                         95
                                               38039
    ge-0/0/4.32767
     Input :
                                                   0
                                                                0
                          0
      Output:
                         95
                                               38039
  Marker Statistics: Marker Rx
                                               Unknown Rx
                                                           Illegal Rx
                                     Resp Tx
    ge-0/0/3.32767
                               0
                                           0
                                           0
    ge-0/0/4.32767
                               0
  Protocol multiservice, MTU: Unlimited, Generation: 191, Route table: 0
    Flags: None
    Policer: Input: __default_arp_policer__
```

SEE ALSO

distribution-list

targeted-distribution

MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet Interfaces

IN THIS SECTION

Benefits | 91

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 based on the command 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.

What Is Enhanced LAG?

IN THIS SECTION

Benefits | 92

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 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.

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
19.3	In Junos OS Release 19.3 and later, for MPC10E and MPC11E MPCs, you cannot apply firewall filters on the micro-BFD packets received on the aggregated Ethernet interface. For MPC1E through MPC9E, you can apply firewall filters on the micro-BFD packets received on the aggregated Ethernet interface only if the aggregated Ethernet interface is configured as an untagged interface.
16.1	Beginning with Junos OS Release 16.1, you can also configure this feature on MX Series routers with aggregated Ethernet interface address of the remote destination as the neighbor address.
16.1	Beginning with Release 16.1R2, Junos OS checks and validates the configured micro-BFD local-address against the interface or loopback IP address before the configuration commit.

14.1	Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases before Junos OS Release 16.1, you must configure the loopback address of the remote destination as the neighbor address.
13.3	Starting with Junos OS Release 13.3, IANA has allocated 01-00-5E-90-00-01 as the dedicated MAC address for micro BFD.

RELATED DOCUMENTATION

Circuit and Translational Cross-Connects Overview

Link Protection of Aggregated Ethernet Interfaces

IN THIS SECTION

- Configuring Aggregated Ethernet Link Protection | 93
- Configuring Aggregated Ethernet Minimum Links | 96
- Example: Configuring Aggregated Ethernet Link Protection | 97

This topic provides information about how to provide link protection for aggregated Ethernet Interfaces and how to configure the minimum number of links in an aggregated Ethernet interfaces bundle.

Configuring Aggregated Ethernet Link Protection

IN THIS SECTION

- Onfiguring Link Protection for Aggregated Ethernet Interfaces | 94
- Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces | 94
- Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link | 95
- Disabling Link Protection for Aggregated Ethernet Interfaces | 96

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

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

NOTE: Link protection is not supported on MX80.

Configuring Link Protection for Aggregated Ethernet Interfaces

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

SEE ALSO

link-protection

aggregated-ether-options

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

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

To configure a primary link and a backup link:

1. Configure the primary logical interface.

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

2. Configure the backup logical interface.

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

SEE ALSO

802.3ad

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link

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

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

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

SEE ALSO

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

Disabling 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)

Configuring Aggregated Ethernet Minimum Links

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

To configure the minimum number of links:

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

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

2. Configure the minimum number of links.

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

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for minimum-links *number* is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links *number* is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

On EX8200 switches, the range of valid values for minimum-links *number* is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

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

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

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

SEE ALSO

aggregated-ether-options minimum-links

Example: Configuring 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;
```

```
}
```

SEE ALSO

aggregated-ether-options

RELATED DOCUMENTATION

Load Balancing on Aggregated Ethernet Interfaces

Performance Monitoring on Aggregated Ethernet Interfaces | 141

Scheduling on Aggregated Ethernet Interfaces

IN THIS SECTION

- Configuring Shared Scheduling on Aggregated Ethernet Interfaces | 98
- Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection | 99

You can configure shared scheduling on aggregated Ethernet Interfaces in link-protection mode or without link protection. The following topic describes how to configure shared scheduling on aggregated Ethernet Interfaces.

Configuring 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 M320 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 aex aggregated-ether-options]
user@host# top
[edit]
user@host# edit interfaces aex shared-scheduler
```

SEE ALSO

aggregated-ether-options

link-protection

shared-scheduler

Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection

On aggregated Ethernet interfaces, you can configure scheduler in non-link-protect mode on the following platforms:

- MX-Series
- M120 and M320 with IQ2 PIC
- T-series platforms (T620 and T320) with IQ2 PIC

The scheduler functions supported are:

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

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

[edit interfaces aeX hierarchical-scheduler]

Prior to Junos OS Release 9.6, the hierarchical scheduler mode on these models required the aggregated-ether-options statement link-protection option. If a link-protection option is not specified, the scheduler is configured in non-link-protect mode.

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 ae X] hierarchy level. The default setting is scale.

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

NOTE: 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.

SEE ALSO

Configuring Hierarchical CoS for a Subscriber Interface of Aggregated Ethernet Links

Junos OS Class of Service User Guide for Routing Devices

RELATED DOCUMENTATION

Understanding Ethernet Link Aggregation on ACX Series Routers | 148

Performance Monitoring on Aggregated Ethernet Interfaces | 141

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 AggregationOverview | 101

- Understanding Aggregated Ethernet LoadBalancing | 102
- Stateful Load Balancing for Aggregated
 Ethernet Interfaces Using 5-Tuple Data | 104
- Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | 107
- Configuring Adaptive Load Balancing | 108
- Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers | 110
- Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers | 117
- Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers | 120
- Example: Configuring Aggregated Ethernet
 Load Balancing | 123

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.

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 128 LAG bundles on M Series, and T Series routers, and 480 LAG bundles on MX Series routers and EX9200 switches. 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 Layer 2 payload by default. When the hash-mode is set to Layer 2 payload, the hashing algorithm uses the IPv4 and IPv6 payload fields for hashing. 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. 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 a Layer 2 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.

NOTE: ACX7100 ae member interfaces do not load balance egress traffic.

• 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.

NOTE: Adaptive load balancing is not supported if the VLAN ID is configured on the aggregated Ethernet interface. This limitation affects the PTX Series Packet Transport Routers and QFX10000 switches only.

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.

NOTE: The pps and scan-interval optional keywords are supported on PTX Series Packet Transport Routers only.

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 Aggegated Ethernet Interfaces or LAG
 Bundles | 106

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. In Junos OS releases earlier than Release 13.2R1, on MX Series routers with Trio-based FPCs (MPCs), the selection of a member link of the ae interface bundle or the next-hop (or unilist of next-hops) for equal-cost multipath ECM) links is performed using a balanced mode next-hop selection methodology and an unbalanced mode of member link or next-hop selection methodology. 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. Starting with Junos OS Release 13.2R1, the capability to perform uniform load balancing and also perform rebalancing is introduced on MX Series routers with MPCs, except MPC3Es and MPC4Es. Rebalancing is not supported when load-balancing is skewed or distorted owing to a change in the number of flows.

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 at the [edit interfaces ae% 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 ae% 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 ae% 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 recreates them newly.

Guidelines for Configuring Stateful Load Balancing for Aggegated 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, MPC3E-3D, MPC5E, and MPC6E line cards. Any other type of line card does not

rigger 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 ae% 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 ae X 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 ae X 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 ae X 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 aeX 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
```

3. 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
```

4. (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

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 | 110
- Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers | 110
- Configuring Symmetrical Load Balancing on Trio-Based MPCs | 114
- Example Configurations | 116

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.

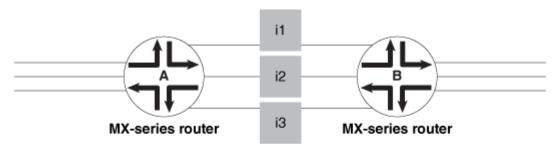
Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

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.

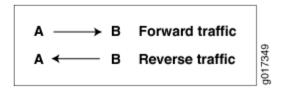
These operations can also be performed at the PIC level by specifying a *hash key*. To configure a hash key at the PIC level, use the symmetric-hash or symmetric-hash complement statement at the [edit chassis hash-key family inet] and [edit chassis hash-key family multiservice] hierarchy levels.

Consider the example in Figure 2 on page 111.

Figure 2: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers



Intelligent Devices



Router A is configured with symmetric hash and Router B is configured with symmetric hash complement. Thus, for a given flow fx, post hash computation is from Router A to Router B through i2. The reverse traffic for the same flow fx 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 {
```

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.

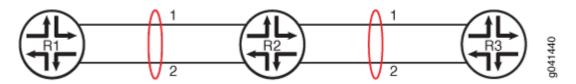
To exclude source and destination MAC addresses from the enhanced-hash-key computation, include the no-mac-addresses statement at the [edit forwarding-options enhanced-hash-key family multiservice] hierarchy level. This statement is disabled by default.

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 3 on page 115, 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 3: Traffic Polarization on Cascaded Routers When Symmetrical Load Balancing in 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 | 116
- Example Configurations of Per-Packet-Forwarding-Engine Settings | 117

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

family		
hash-key		
inet		
multiservice		
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 | 120
- Configuring Symmetrical Hashing for family inet on Both Routers | 121
- Configuring Symmetrical Hashing for family inet and family multiservice on the Two Routers | 122

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 | 123

Example: Configuring Aggregated Ethernet Load Balancing

IN THIS SECTION

- Requirements | 123
- Overview | 123
- Configuration | 126
- Verification | 139

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 | 125

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: Adaptive load balancing is not supported if the VLAN ID is configured on the aggregated Ethernet interface. This limitation affects the PTX Series Packet Transport Routers only.

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.

NOTE: The pps and scan-interval optional keywords are supported on PTX Series Packet Transport Routers only.

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.

ae0 120.168.104.1/30 ge-1/3/0 ge-5/3/0 ge-1/3/1 ge-5/3/1 ge-1/3/2 ge-5/3/2 ge-1/3/3 ge-5/3/3 ge-1/3/4 ge-5/3/4 ge-1/0/0 ge-1/2/0 ael 120.168.100.2/30 120.168.100.1/30 120.168.105.1/30 ge-1/0/1 ge-2/2/1 ge-5/0/1 120.168.101.2/30 120.168.101.1/30 20.168.0.2/32 120.168.0.4/32 120.168.0.9 ge-2/2/2 ge-5/0/2 ge-2/2/3 ge-5/0/3 ge-2/2/4 ge-5/0/4 ge-2/2/5 ge-5/0/5 ge-2/2/6 ge-5/0/6 ge-2/2/7 ge-5/0/7 ge-5/0/8 ge-2/2/8

Figure 4: Aggregated Ethernet Load Balancing

Configuration

IN THIS SECTION

- CLI Quick Configuration | 126
- Configuring Adaptive Load Balancing | 132
- Results | **135**

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 <code>[edit]</code> hierarchy level.

```
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-1/2/1 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
```

```
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 aeO 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-2/2/6 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 aeO aggregated Ethernet bundle.

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

7. Configure LACP on the aeO aggregated Ethernet bundle.

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

8. Configure the interface parameters for the aeO 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 | 140

Confirm that the configuration is working properly.

Verifying Adaptive Load Balancing on ae0

Purpose

Verify that packets received on the aeO 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
                                                                 bps
                                                  Bytes
    Bundle:
        Input :
                       848761
                                               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
                                               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:
                                               13374435
                                                                1280
                       146678
                                       1
        Output:
                  33205071207
                                  697870 25374651121458
                                                           4269487384
```

Meaning

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

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.	
13.2R1	Starting with Junos OS Release 13.2R1, the capability to perform uniform load balancing and also perform rebalancing is introduced on MX Series routers with MPCs, except MPC3Es and MPC4Es.	
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

IN THIS SECTION

- ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces Overview | 142
- Guidelines for Configuring Performance Monitoring Functionalities on Aggregated Ethernet Interfaces | 144

Use this topic to understand ore about performance monitoring features on aggregated Ethernet Interfaces. You can refer to the guidelines for configuring performance monitoring features before you configure performance monitoring.

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

Starting with Junos OS Relase 16.1R1, you can 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). These performance monitoring functionalities are supported on the following platforms:

- MX Series routers with 16-port 10-Gigabit Ethernet MPCs and Trio-based FPCs (MPCs), where the same level of support for the Ethernet services OAM mechanisms on non-aggregated Ethernet interfaces is available on AE interfaces
- MX2020 routers
- ETH-DM is supported on MPC3E and MPC4E modules with only software timestamping
- ETH-SLM is supported on MPC3E and MPC4E modules.

Also, 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 gracful Routing Engine switchover (GRES) for CFM sessions.

Connectivity fault management (CFM) sessions operate in centralized mode over AE interfaces by default. Y.1731 performance monitoring (PM) is supported on centralized CFM sessions over AE interfaces. Also, distribution of CFM session over AE interfaces to line cards is supported from Junos OS Release 13.3. 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. For software timestamping, ETH-DM PDUs need to be transmitted and received on the same line card (same member of the AE interface). All the received ETH-DM PDUs are always redirected to the anchor Packet Forwarding Engine. In the transmission path, if the interface on the anchor Packet Forwarding Engine goes down, then the OAM pdus are redirected to one of the subordinate or member FPCs. Therefore, the processing of ETH-DM PDUs always occurs at the CPU of the line card or module that hosts the

anchor Packet Forwarding Engine. ETH-DM is supported on AE interfaces with CCC, bridge, virtual private LAN service (VPLS), and inet address families. ETH-DM is supported for both active-active and active-standby modes of AE interfaces. For one-way delay measurement (1DM), the system clocks of the initiator MEP that transmits a request frame and the responder MEP that receives a reply frame need to be synchronized.

For Ethernet loss measurement on AE interfaces, with the active-standby mode of the interfaces, transmission and reception of PDUs is always through the Packet Forwarding Engine that hosts the active link. For the active-standby mode of the AE interfaces, you can configure a maximum of only two member links. ETH-LM is supported only when all the active member or child links are on the same Packet Forwarding Engine. For the downstream maintenance endpoints (MEPs), ETH-LM is supported for CCC, VPLS, and bridge address families, and for upward MEPs, ETH-LM is supported only for CCC families. In the transmission path, with active-standby links of AE interfaces, whenever the active child link fails, if the standby link is non-local, the packets are redirected to the new active link. When this redirection occurs, the ETH-LM counters are reset. If the standby link is on same Packet Forwarding Engine as the active link, then the counters are not reset because the counters are read on the local Packet Forwarding Engine memory and to prevent the other end of the session to treat new Packet Forwarding Engine counters as losses owing to reset of the counters. In the received path, with activestandby links of AE interfaces, all the child links are programmed in the input list using next-hops to redirect the packets to the anchor FPC after copying the counters in the Packet Forwarding Engine. For Ethernet synthetic loss measurement (SLM), processing of SLM PDUs for requests and responses similar to other protocols from the line card CPU is implemented. All other computation and data are softwarebased. ETH-SLM is supported on AE interfaces for CCC, bridge, VPLS, and inet families.

NOTE: Starting with Junos OS Release 16.1, 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. Starting with Junos OS Release 16.2, connectivity fault management (CFM) is supported when enhanced LAG is enabled.

Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the performance-monitoring statement and its substatements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level) is not supported when the network-to-network (NNI) or egress interface is an aggregated Ethernet interface with member links on DPCs.

Before you can start an ETH-DM, ETH-LM, or ETH-SLM measurement sessions across an aggregated Ethernet service, you must configure two MX Series 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.

SEE ALSO

ITU-T Y.1731 Ethernet Service OAM Overview

Guidelines for Configuring Performance Monitoring Functionalities 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.
- 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.
- For ETH-LM over AE interfaces, in an active-standby setup, whenever the link failover from the active interface to the standby interface happens, the counters are reset.

SEE ALSO

Configuring Ethernet Frame Delay Measurement Sessions

Configuring Ethernet Frame Loss Measurement

Configuring Ethernet Synthetic Loss Measurements

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
16.2	Starting with Junos OS Release 16.2, connectivity fault management (CFM) is supported when enhanced LAG is enabled.
16.1R1	Starting with Junos OS Relase 16.1R1, you can 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.
16.1	Starting with Junos OS Release 16.1, Ethernet loss measurement over an aggregated Ethernet (ae) interface is not supported when the enhanced LAG functionality is enabled on a router.
16.1	Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the performance-monitoring statement and its substatements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level) is not supported when the network-to-network (NNI) or egress interface is an aggregated Ethernet interface with member links on DPCs.

RELATED DOCUMENTATION

Aggregated Ethernet Interfaces | 63

ITU-T Y.1731 Ethernet Service OAM Overview

Periodic Packet Management

IN THIS SECTION

Understanding Periodic Packet Management on MX Series Routers | 146

Configuring Periodic Packet Management on MX Series Routers | 147

Periodic packet management (PPM) 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.

Understanding Periodic Packet Management on MX Series Routers

Periodic packet management (PPM) for MX Series routers is responsible for processing a variety of timesensitive 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.

BEST PRACTICE: 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.

Configuring Periodic Packet Management on MX Series Routers

IN THIS SECTION

Enabling Centralized Periodic Packet Management | 147

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.

Enabling 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 PFE failure.

BEST PRACTICE: 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.

To enable centralized periodic packet management:

1. From configuration mode, enable centralized periodic packet management by specifying the nodelegate-processing statement at the [edit] hierarchy level.

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

2. Commit the configuration by using the commit statement.

[edit]
user@host# commit

3. 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
```

4. 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.

RELATED DOCUMENTATION

Ensuring That Distributed ppm Is Not Disabled

Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)

ppm

Understanding Ethernet Link Aggregation on ACX Series Routers

IN THIS SECTION

- Load Balancing | 151
- LACP Monitoring | 152
- Link Protection | 153
- Understanding the Algorithm Used to Hash LAG Bundle | 154

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.

NOTE: ACX Series routers support connectivity fault management (CFM) on aggregated Ethernet interfaces with continuity check interval of 100 milliseconds or higher.

NOTE: ACX5048 and ACX5096 routers support connectivity fault management (CFM) on aggregated Ethernet interfaces with continuity check interval of 1 second or higher.

NOTE: The Ethernet options configurations for ACX5048 and ACX5096 routers differ compared to other ACX Series routers. For more information, see Layer 2 Next Generation Mode for ACX Series.

On ACX Series routers, up to 128 AE interfaces can be created with each AE interface having up to 8 physical interfaces. AE interfaces can be created across PICs and fixed-ports on the chassis.

NOTE: On ACX5048 and ACX5096 routers, up to 64 AE interfaces can be created with each AE interface having up to 16 physical interfaces.

ACX Series routers do not support statistics for aggregated Ethernet interface. However, statistics can be retrieved for member interface.

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 (aex), that is, the defined bundle, to be labeled "up":

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

```
[edit interfaces]
user@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.

NOTE: 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 ACX Series 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.

ACX Series 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;
        }
    }
}
```

NOTE:

- 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;
}
```

NOTE: 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;
}
```

Link Protection

Link protection can be configured on AE interfaces to provide 1:1 link resiliency using LACP. Primary and backup links can be configured within an AE bundle. The primary link is used for all transit traffic and host generated traffic. The backup link is used when the primary link fails.

Link protection is supported only when the AE bundles have no more than 2 member links, one primary and another backup. LACP works in revertive link-protection mode by default and can be configured to work in non-revertive mode.

NOTE: Link protection without LACP (static link protection on AE interfaces) is not supported on all ACX Series routers. Link protection works as expected with LACP configured on the AE bundle.

Configuring Link Protection for Aggregated Ethernet Interfaces

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

To configure link protection:

1. 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
```

Disabling Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the **delete interface revert aex** configuration command.

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

Understanding the Algorithm Used to Hash LAG Bundle

ACX Series 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 15 on page 155 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 15: 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
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 16 on page 156 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 16: 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

CoS on ACX Series Routers Features Overview

Controlling Network Access Using Traffic Policing Overview

Overview of Firewall Filter Match Conditions and Actions on ACX Series Routers



Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Interfaces | 158

Configuring Gigabit Ethernet Interfaces

IN THIS CHAPTER

- Configuring 10-Gigabit Ethernet PICs | 158
- Configuring 40-Gigabit Ethernet PICs | 200
- Configuring 100-Gigabit Ethernet MICs/PICs | 203
- Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | 231
- Configuring Layer 2 Overhead Attribute in Interface Statistics | 247
- Configuring Gigabit Ethernet Policers | 255
- Gigabit Ethernet Autonegotiation | 271

Configuring 10-Gigabit Ethernet PICs

IN THIS SECTION

- 10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 159
- 12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 163
- 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 166
- P2-10G-40G-QSFPP PIC Overview | 167
- Configuring the P2-10G-40G-QSFPP PIC | 178
- Example: Configuring the P2-10G-40G-QSFPP PIC | 182
- Framing Overview | 187
- Understanding WAN Framing | 188
- Configuring Ethernet Framing | 188
- Modes of Operation | 189
- Configuring Mixed-Rate Mode Operation | 190
- Configuring Line-Rate Mode on PICs Supporting Oversubscription | 192

- Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC | 192
- Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 194
- Gigabit Ethernet Notification of Link Down Alarm Overview | 197
- Notification of Link Down for Optics Options Overview | 198
- Configuring Gigabit Ethernet Notification of Link Down Alarm | 198
- Configuring Link Down Notification for Optics Options Alarm or Warning | 198

You can learn about the 10-Gigabit Ethernet PICs in this topic. You can configure LAN and WAN framing, modes of operation, and alarm generation when link is down.

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview

This section describes the main features and caveats of the 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP) and specifies which routers support this PIC.

The 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) is supported on Juniper Networks T640 Core Routers, T1600 Core Routers, and T4000 Core Routers. It has the following features:

- Access to all 10-Gigabit Ethernet port counters through SNMP
- Intelligent handling of oversubscribed traffic in applications such as data centers and dense-core uplinks
- Line-rate operation for five 10-Gigabit Ethernet ports from each port group, or a total WAN bandwidth of 100 Gbps with Packet Forwarding Engine bandwidth of 50 Gbps
- Flexible encapsulation, source address and destination address media access control (MAC) filtering, source address MAC learning, MAC accounting, and MAC policing
- Interface encapsulations, such as the following:
 - ethernet-ccc—Ethernet cross-connect
 - vlan-ccc—802.1Q tagging for a cross-connect
 - ethernet-tcc—Ethernet translational cross-connect
 - vlan-tcc—Virtual LAN (VLAN) translational cross-connect
 - extended-vlan-ccc—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect
 - ethernet-vpls—Ethernet virtual private LAN service

- vlan-vpls—VLAN virtual private LAN service
- flexible-ethernet-services—Allows per-unit Ethernet encapsulation configuration
- WAN PHY features, such as the following:
 - WAN PHY mode on a per-port basis
 - Insertion and detection of path trace messages
 - Ethernet WAN Interface Sublayer (WIS) object

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).

- Single, stacked, and flexible VLAN tagging modes
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames
- Link aggregation group (LAG) on single chassis
- Interoperability with other 10-Gigabit Ethernet PICs in M Series and T Series routers in the LAN PHY and WAN PHY modes
- Interrupt-driven link-down detection mechanism
- Two-to-one oversubscription of traffic across a port group

Traffic from 10 ingress ports to the Packet Forwarding Engine traffic is statically mapped to one of the 5 egress ports. 10 Gbps of bandwidth toward the Packet Forwarding Engine is shared by two ingress ports (called a *port group*), thereby achieving two-to-one oversubscription. This scheme provides two-to-one oversubscription across a port group and not across the entire PIC.

- Four queues per physical interface on ingress and eight queues per physical interface on egress
- A separate control queue per physical interface to ensure that the control packets are not dropped during oversubscribed traffic. The control queue can be disabled in the CLI.
- Optical diagnostics
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification

- Weighted round-robin scheduling with two queue priorities (low and strict-high)
- Committed information rate and peak information rate shaping on a per-queue basis
- Excess information rate configuration for allocation of excess bandwidth
- IEEE 802.3ah Operation, Administration, and Maintenance (OAM)-related operations, such as the following:
 - Link fault management
 - Link discovery
 - Graceful Routing Engine Switchover
- IEEE 802.3ag Operation, Administration, and Maintenance (OAM)-related operations, such as the following:
 - Connectivity fault management (CFM)
 - Linktrace
 - Loopback
 - Graceful Routing Engine switchover (GRES)

The 10-port 10-Gigabit Ethernet LAN/WAN PIC has the following caveats:

- Source address and destination address MAC filtering takes place after oversubscription is handled.
- Oversubscription on the PIC operates across a port group of two ports and not at the PIC level.
- Queuing is not supported at the *logical interface* level.
- Committed information rate and peak information rate configurations are not supported at the physical interface level.
- There is limited packet buffering of 2 MB.
- Delay-bandwidth buffering configuration is not supported.
- Multifield classifiers are not supported at the PIC level.

The multifield classification can be done at the Packet Forwarding Engine using the firewall filters, which overrides the classification done at the PIC level. The multifield classification at the Packet Forwarding Engine occurs after the PIC handles the oversubscribed traffic.

- Egress MAC policer statistics not supported.
- Byte counters are not supported at the gueue level.

- Only TPID (0x8100) is supported.
- Line-timing mode is not supported.
- MAC-level Rx VLAN tagged frames counter is not supported.
- OAM unified in-service software upgrade (unified ISSU) is not supported.
- OAM remote loopback is not supported.

The 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) supports link aggregation. For bandwidth aggregation, load sharing, and link protection, LAG can be enabled. Once aggregated Ethernet is enabled, Link Aggregation Control Protocol (LACP) forms an aggregated bundle of member links.

Only features that are supported across all of the linked devices will be supported in the resulting LAG bundle. The following caveats apply to LAG bundles that involve 10-port 10-Gigabit Ethernet LAN/WAN PIC (PD-5-10XGE-SFPP) ports:

- Non-standard TPID for VLAN tagging is not supported, except for 0x8100.
- The number of user created IFLs is limited to 4065/PIC and 1022/port.
- Classifier tables are limited to 8 for each BA classifier type.
- Forwarding classes are limited to 8.
- The guaranteed-rate and shaping-rate statements are not supported at the IFD level.
- The per-unit-scheduler and hierarchical-scheduler statements are not supported.
- Only the strict-high and low levels of scheduling priorities are supported.
- The excess-priority configuration is not supported.
- The buffer-size configuration under schedulers is not supported.
- WRED is not supported.
- srTCM and trTCM are not supported.
- Shared scheduler mode is not supported.

Table 17 on page 163 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP).

Table 17: Capabilities of 10-Gigabit Ethernet LAN/WAN PICs

Capability	Support
Maximum VLANs per PIC	4065
Maximum VLANs per port	1022
MAC learning per port	960
MAC accounting per port	960
MAC filtering per port	960 (64 filters per physical or logical interface) 960 filters across multiple logical interfaces
MAC policers	128 ingress Mac policers 128 egress Mac policers
Classifers	Eight classifiers per PIC for each BA classifier type

SEE ALSO

Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 194

IEEE 802.1ag OAM Connectivity Fault Management Overview

IEEE 802.3ah OAM Link-Fault Management Overview

12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC is a 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number, PF-12XGE-SFPP) on T4000 Core Routers.

The following features are supported on the 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC:

• Access to all 10-Gigabit Ethernet port counters through SNMP.

- Logical interface—level MAC filtering, accounting, policing, and learning for source media access control (MAC).
- Flexible encapsulation.
- Single, stacked, and flexible VLAN tagging modes.
- Native VLAN configuration to allow untagged frames to be received on the tagged interfaces.
- Maximum transmission unit (MTU) size of up to 9192 bytes for Ethernet frames.
- Link aggregation group (LAG) on single chassis.
- Interoperability with other 10-Gigabit Ethernet PICs on M Series and T Series routers in LAN PHY mode.
- Eight queues per physical interface on egress.
- Behavior aggregate (BA) classification (IPv4 DSCP, IPv6 DSCP, Inet precedence, IEEE 802.1P, IEEE 802.1AD, MPLS EXP) and fixed classification.
- Defining the VLAN rewrite operation to be applied to the incoming and outgoing frames on logical interfaces on this PIC.

NOTE: Only the Tag Protocol Identifier (TPID) 0x8100 is supported.

- Interface encapsulations, such as the following:
 - untagged—Default encapsulation, when other encapsulation is not configured.
 - You can configure only one logical interface (unit 0) on the port.
 - You cannot include the vlan-id statement in the configuration of the logical interface.
 - vlan-tagging—Enable VLAN tagging for all logical interfaces on the physical interface.
 - stacked-vlan-tagging—Enable stacked VLAN tagging for all logical interfaces on the physical interface.
 - ethernet-ccc—Ethernet cross-connect.
 - ethernet-tcc—Ethernet translational cross-connect.
 - vlan-ccc—802.1Q tagging for a cross-connect.
 - vlan-tcc—Virtual LAN (VLAN) translational cross-connect.
 - extended-vlan-ccc—Standard Tag Protocol Identifier (TPID) tagging for a cross-connect.

- extended-vlan-tcc—Standard Tag Protocol Identifier (TPID) tagging for an Ethernet translational cross-connect.
- ethernet-vpls—Ethernet virtual private LAN service.
- vlan-vpls—VLAN virtual private LAN service.
- flexible-ethernet-services—Allows per-unit Ethernet encapsulation configuration.
- The following Layer 3 protocols are also supported:
 - IPv4
 - IPv6
 - MPLS
- WAN PHY features, such as the following:
 - WAN PHY mode on a per-port basis.
 - Insertion and detection of path trace messages.
 - Ethernet WAN Interface Sublayer (WIS) object.

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on T4000 routers with 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+.

The 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC does not support:

• MAC filtering, accounting, and policing for destination MAC at the logical interface level.

NOTE: Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This enables the OSPF protocol to work.

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.
- Multiple TPIDs

Capability	Support
Maximum logical interfaces per PIC	32,000
Maximum logical interfaces per port	For IPv4 the limit is 4093. For IPv6 the limit is 1022.
Classifers	Eight classifiers per PIC for each BA classifier type

SEE ALSO

Configuring 40-Gigabit Ethernet PICs | 200

24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview

This section describes the main features and caveats of the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP).

The following major software features are supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP):

- Twenty-four 10-Gigabit Ethernet interfaces in two-to-one oversubscription of traffic in oversubscribed mode or 12 ports in line-rate mode. For more information about oversubscribed mode and line-rate mode, see the "Configuring Line-Rate Mode on PICs Supporting Oversubscription" on page 192.
- Traffic is classified as control traffic or best-effort traffic with non-class-of-service-aware tail drops of best-effort traffic in oversubscribed mode.

The aggregate bandwidth of all the ports together is 120 Gbps. No hard partitioning of bandwidth is done—that is, if one port group is active, it can support 120 Gbps traffic. The bandwidth for best-effort traffic is shared among all the 24 ports.

Note that the preclassification is restricted to two traffic classes, and is not user-configurable.

- All Junos OS configuration commands supported on the existing 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The output of the show interfaces extensive *operational mode command* now displays preclassification queue counters.

- Line-rate mode operation of the first 12 ports can be achieved by using the [set chassis fpc *fpc-number* pic *pic-number* linerate-mode] command. By default, the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ works in oversubscribed mode.
- LAN PHY mode and WAN PHY mode on a per-port basis. WAN PHY mode can be achieved by using the [set interfaces *interface-name* framing wan-phy] command.
- WAN PHY features, such as the following:
 - Insertion and detection of path trace messages.
 - Ethernet WAN Interface Sublayer (WIS) object.
- Aggregated Ethernet is supported only in line-rate mode.
- Link aggregation group (LAG) is supported only in line-rate mode.
- 4000 logical interfaces per physical interface and 32,000 logical interfaces per chassis.
- Access to all 10-Gigabit Ethernet port counters through SNMP.

NOTE: *Graceful Routing Engine switchover* (GRES) and *nonstop active routing* (NSR) are now supported on T4000 routers.

SEE ALSO

Configuring Gigabit Ethernet Policers | 255

Gigabit Ethernet Autonegotiation | 271

P2-10G-40G-QSFPP PIC Overview

IN THIS SECTION

- Understanding Dual Configuration on P2-10G-40G-QSFPP PIC | 168
- Understanding Port Group | 169
- Port Numbering on P2-10G-40G-QSFPP PIC When Port Groups Are Not Configured | 173
- 10-Gigabit Ethernet Mode | 176
- 40-Gigabit Ethernet Mode | 178

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

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.

Starting from Junos OS Release 14.2R3 and 16.1R1, 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.

NOTE: 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.

NOTE: If you want 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 on PTX5000 routers 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 18 on page 170 shows the port numbering in 40-Gigabit Ethernet mode and in 10-Gigabit Ethernet mode at the port group level.

Table 18: 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
O (O)	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 18: 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-1/1/1	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
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

Table 18: 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
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
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 18: 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
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
	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 19 on page 174 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 19: 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
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

Table 19: 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
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
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

Table 19: 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
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
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.

NOTE: 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* otnoptions preemptive-fast-reroute] hierarchy level.
- Configure thresholds and interval for the optical channel data unit (ODU) signal degradation (odusignal-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.

SEE ALSO

Configuring 100-Gigabit Ethernet MICs/PICs | 203

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

IN THIS SECTION

- Configuring the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 178
- Configuring the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 179
- Configuring the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 179
- Configuring the PIC at Port Group Level | 180
- Configuring Framing Mode on P2-10G-40G-QSFPP PIC | 181

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.

Configuring 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)
```

Configuring 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.

Configuring 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.

Configuring 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)
```

NOTE: 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.

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

IN THIS SECTION

- Configuring LAN PHY or WAN PHY Framing Mode in 10-Gigabit Ethernet Mode | 181
- Configuring LAN PHY Framing Mode in 40-Gigabit Ethernet Mode | 182

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:

Configuring 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:

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.

Configuring 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.

RELATED DOCUMENTATION

```
P2-10G-40G-QSFPP PIC Overview | 167
```

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

IN THIS SECTION

- Requirements | 183
- Overview | 183
- Configuration | 183

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 | 185

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
```

2. 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:

1. 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 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.

1. 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 | 186

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
   Errored blocks
 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.

SEE ALSO

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, T640, T1600, T4000 and PTX Series Packet Transport routers only.

NOTE: The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).

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. For information about SONET options supported on this interface, see Configuring SONET Options for 10-Gigabit Ethernet Interfaces.

NOTE: SONET or SDH framing mode configuration framing (sdh | sonet) is not applicable on the 10-Gigabit Ethernet ports. Configuring the wan-phy framing mode on the 10-Gigabit Ethernet ports allows the interface to accept SONET or SDH frames without further configuration.

SEE ALSO

Configuring SONET/SDH Framing Mode for Ports

Configuring 100-Gigabit Ethernet MICs/PICs | 203

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 in 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.

SEE ALSO

framing

Configuring Ethernet Framing | 188

Framing Overview | 187

Ethernet Interfaces User Guide for Routing Devices

Configuring Ethernet Framing

The 10-Gigabit Ethernet interfaces uses the interface type xe-fpcl picl 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);
```

NOTE:

- The T4000 Core Router supports only LAN PHY mode in Junos OS Release 12.1R1. Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP). Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
- 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.

NOTE:

- SONET or SDH framing mode configuration framing (sdh | sonet) is not applicable on the 10-Gigabit Ethernet ports. Configuring the wan-phy framing mode on the 10-Gigabit Ethernet ports allows the interface to accept SONET or SDH frames without further configuration.
- If you configure the WAN PHY mode on an aggregated Ethernet interface, you must set the aggregated Ethernet link speed to OC192.

SEE ALSO

Configuring 100-Gigabit Ethernet MICs/PICs | 203

Modes of Operation

10-Gigabit Ethernet PICs operate in the following modes:

- Line-rate mode—By default, the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP) operates in line-rate mode.
 - In a 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP), 12 ports (ports 0–11)can operate in line-rate mode . To configure the PF-24XGE-SFPP PIC to operate in line-rate mode, include the linerate-mode statement at the [edit chassis set fpc *fpc-number* pic *pic-number*] hierarchy level.
- Oversubscribed mode—In this mode, all ports on the PIC are enabled with two-to-one oversubscription. In a PF-24XGE-SFPP PIC, by default, two-to-one oversubscription of traffic is achieved in oversubscribed mode—Traffic from 24 ingress ports to the Packet Forwarding Engine is statically mapped to one of the 12 egress ports. 10 Gbps of bandwidth traffic moving toward the Packet Forwarding Engine is shared by two ingress ports (called a port group), thereby achieving two-to-one oversubscription. This scheme provides two-to-one oversubscription across a port group and not across the entire PIC.

NOTE: PF-12XGE-SFPP PIC always operates at line rate.

• Mixed-rate mode or dual-rate mode—Dual-rate mode or mixed-rate mode for PF-24XGE-SFPP allows you to configure a mix of port speeds of 1 Gbps and 10 Gbps. However, on PF-12XGE-SFPP, note that you can configure port speeds of either 1 Gbps and 10 Gbps when the PIC is in line rate mode. You can enable mixed-rate mode and set port speeds with the mixed-rate-mode and speed 16 [106 statements respectively at the [edit chassis fpc x pic y] hierarchy level. You can disable mixed-rate mode with the delete chassis fpc x pic y mixed-rate-mode statement.

NOTE: To change the port speed from 10 Gbps to 1 Gbps on the PF-24XGE-SFPP and PF-12XGE-SFPP PICs, SFP optics is required.

SEE ALSO

mixed-rate-mode

Configuring Mixed-Rate Mode Operation

To configure mixed-rate mode operation for a PF-24XGE-SFPP PIC:

1. Navigate to the [edit chassis] hierarchy level.

2. On a T4000 router, configure the mixed-rate mode by including the mixed-rate-mode statement at the [edit chassis fpc *slot-number* pic *pic-number*] hierarchy level.

```
[edit chassis]
user@host# set fpc fpc-slot pic pic-number mixed-rate-mode
```

On an LCC in a routing matrix, configure the mixed-rate mode by including the mixed-rate-mode statement at the [edit chassis lcc *lcc number* fpc *slot-number* pic *pic-number*] hierarchy level.

```
[edit chassis]
user@host# set lcc lcc number fpc fpc-slot pic pic-number mixed-rate-mode
```

3. Specify the port and the port speed that need to be configured. You can use one of the following speed attributes for this configuration.

```
[edit chassis]
user@host# set fpc fpc-slot pic pic-number port port-number speed 1G;
user@host# set fpc fpc-slot pic pic-number port port-number speed 10G;
user@host# set lcc lcc number fpc fpc-slot pic pic-number speed 1G;
user@host# set lcc lcc number fpc fpc-slot pic pic-number speed 10G;
```

NOTE: On a 12 port 10-Gigabit Ethernet PIC (PF-12XGE-SFPP), you can configure the port speed as 16 by including the **set** fpc *fpc-slot* pic *pic-number* port port-number speed **1G** statement at the <code>[edit chassis]</code> hierarchy level.

NOTE: To change the port speed from 10 Gbps to 1 Gbps on PF-24XGE-SFPP and PF-12XGE-SFPP PICs, SFP optics is required.

To disable mixed-rate mode operation, include the **delete chassis fpc** x **pic** y **mixed-rate-mode** statement at the [edit chassis] hierarchy level.

SEE ALSO

Modes of Operation | 189

mixed-rate-mode

Configuring Line-Rate Mode on PICs Supporting Oversubscription

For 10-Gigabit Ethernet LAN/WAN PICs supporting oversubscription, oversubscribed Ethernet mode is set by default. To configure these PICs in line-rate mode, include the linerate-mode statement at the [edit chassis set fpc *fpc-number* pic *pic-number*] hierarchy level:

[edit chassis]

set fpc fpc-number pic pic-number linerate-mode;

To return to the default oversubscribed Ethernet mode, delete the linerate-mode statement at the [edit chassis fpc *fpc-number* pic *pic-number*] hierarchy level.

NOTE: When the mode of operation of a PIC is changed, the PIC is taken offline and then brought back online immediately.

The following 10-Gigabit Ethernet LAN/WAN PICs support line-rate mode:

- 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP)
- 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP)

SEE ALSO

Configuring 100-Gigabit Ethernet MICs/PICs | 203

Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC

Table 20 on page 192 lists the scenarios of handling oversubscription on the 10-port 10-Gigabit Ethernet LAN/WAN PIC for different combinations of port groups and active ports on the PIC.

Table 20: Handling Oversubscription on 10-Gigabit Ethernet LAN/WAN PICs

Number of Port Groups with Two Active Ports (A)	Number of Port Groups with One Active Port (B)	Total Number of Ports Used on PIC (C = Ax2 + B)	Status of Oversubscription and Throughput
0	1	1	Oversubscription is not active. Each port will receive 10 Gbps throughput.

Table 20: Handling Oversubscription on 10-Gigabit Ethernet LAN/WAN PICs (Continued)

Number of Port Groups with Two Active Ports (A)	Number of Port Groups with One Active Port (B)	Total Number of Ports Used on PIC (C = Ax2 + B)	Status of Oversubscription and Throughput
0	2	2	Oversubscription is not active. Each port will receive 10 Gbps throughput.
0	5	5	Oversubscription is not active. Each port will receive 10 Gbps throughput.
1	0	2	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
1	4	6	Oversubscription is active for the port group that has two active ports. Each port in this port group will receive 5 Gbps throughput (with default shaper configuration). For the remaining four ports, oversubscription is not active. Each port will receive 10 Gbps throughput.
3	0	6	Oversubscription is active. Each port will receive 5 Gbps throughput (with default shaper configuration).
5	0	10	Oversubscription is active on all 10 ports (5 port groups). Each port will receive 5 Gbps throughput (with default shaper configuration).

SEE ALSO

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 159

Configuring Line-Rate Mode on PICs Supporting Oversubscription | 192

Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 194

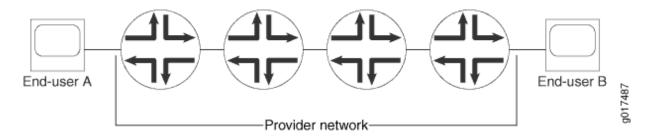
Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC

On a 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), a control queue is used to queue all control packets received on an ingress port. This ensures that control protocol packets do not get dropped randomly when there is congestion due to oversubscription. The following control protocols are supported:

- OSPF
- OSPF3
- VRRP
- IGMP
- RSVP
- PIM
- BGP
- BFD
- LDP
- IS-IS
- RIP
- RIPV6
- LACP
- ARP
- IPv6 NDP
- Connectivity fault management (CFM)
- Link fault management (LFM)

These control packets can either terminate locally or transit through the router. The control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. Hence, if transit control traffic is taking too much bandwidth, then it can cause drops on locally terminating control traffic, as shown in Figure 5 on page 195.

Figure 5: Control Queue Rate Limiter Scenario



If the end users generate a mass of malicious traffic for which the port number is 179 (BGP), the router dispatches that traffic to the ingress control queue. Further, if congestion occurs in this ingress control queue due to this malicious traffic, the provider's network control packets may be affected.

In some applications, this can be perceived as a new vulnerability. To address this concern, you can disable the control queue feature. With the control queue feature disabled, you must take precautions to protect control traffic through other means, such as mapping control packets (using BA classification) to a queue that is marked strict-high or is configured with a high CIR.

You can disable the control queue for all ports on the PIC. To disable the control queue, use the **set chassis fpc** *n* **pic** *n* **no-pre-classifier** command. By default, the no-pre-classifier statement is not configured and the control queue is operational.

Deleting the no-pre-classifier statement re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.

NOTE:

- This functionality is applicable both in OSE and line-rate modes.
- The control queue feature is enabled by default in both OSE and line-rate modes, which can be overridden by the user configuration.
- When the control queue is disabled, various show queue commands will show *control queue* in the output. However, all control queue counters are reported as zeros.
- Changing this configuration (enabling or disabling the control queue feature) results in the PIC being taken offline and brought back online.

Once the control queue is disabled, the Layer 2/Layer 3 control packets are subject to queue selection based on BA classification. However, some control protocol packets will not be classified using BA classification, because they might not have a VLAN, MPLS, or IP header. These are:

- Untagged ARP packets
- Untagged Layer 2 control packets such as LACP or Ethernet OAM

Untagged IS-IS packets

When the control queue feature is disabled, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to the restricted queue corresponding to the forwarding class associated with queue 0, as shown in the following two examples.

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the forwarding class (FC) associated with queue 0 is "be" (based on the forwarding-class statement configuration). "be" maps to restricted-queue number 3 (based on the "restricted-queue" configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets will go to ingress queue 3 (not to ingress queue 0).

```
[edit chassis]
forwarding-classes {
    queue 0 be;
    queue 1 af-low8;
    queue 2 af-high;
    queue 3 ef;
    queue 4 ops_control;
    queue 5 net_control;
    queue 6 af-low10_12;
   }
    restricted-queues {
        forwarding-class ef queue-num 0;
        forwarding-class af-low8 queue-num 1;
        forwarding-class af-low10_12 queue-num 1;
        forwarding-class af-high queue-num 2;
        forwarding-class be queue-num 3;
   }
```

Forwarding Untagged Layer2 Control Packets to Queue 3

With this configuration, the FC associated with queue 0 is "ef" (based on the forwarding-class statement configuration). "ef" maps to restricted-queue number 0 (based on the restricted-queue statement configuration). Hence, with this particular configuration, untagged ARP, IS-IS, and other untagged Layer 2 control packets would go to ingress queue 0.

For tagged ARP, IS-IS, or Layer2 control packets, users should configure an explicit dot1p/dot1ad classifier to make sure these packets are directed to the correct queue. Without an explicit dot1p/

dot1ad classifier, tagged ARP, IS-IS, or Layer 2 control packets will go to the restricted-queue corresponding to the forwarding class associated with queue 0.

```
Fedit chassis
forwarding-classes {
    queue 0 ef;
                              <<< ef and be are interchanged
    queue 1 af-low8;
    queue 2 af-high;
    queue 3 be;
                              <<< ef and be are interchanged
   queue 4 ops_control;
   queue 5 net_control;
   queue 6 af-low10_12;
    restricted-queues {
        forwarding-class ef queue-num 0;
        forwarding-class af-low8 queue-num 1;
        forwarding-class af-low10_12 queue-num 1;
        forwarding-class af-high queue-num 2;
        forwarding-class be queue-num 3;
   }
```

SEE ALSO

no-pre-classifier

Gigabit Ethernet Notification of Link Down Alarm Overview

Notification of link down alarm generation and transfer is supported for all 10-Gigabit Ethernet PIC interfaces on M120 and M320 routers. On the MX Series and T series routers, notification of link down alarm generation and transfer is supported for all Gigabit Ethernet Interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit).

SEE ALSO

asynchronous-notification

Notification of Link Down for Optics Options Overview

Notification of link down is supported for IQ2 10-Gigabit Ethernet interfaces and MX Series DPCs. You can use link down notification to help identify optical link connectivity problems.

For information on configuring link down notification, see "Configuring Link Down Notification for Optics Options Alarm or Warning" on page 198.

SEE ALSO

Configuring Link Down Notification for Optics Options Alarm or Warning | 198

Configuring Gigabit Ethernet Notification of Link Down Alarm

Notification of link down alarm generation and transfer is supported for all 10-Gigabit Ethernet PIC interfaces on M120 and M320 routers. On the MX Series and T Series routers, notification of link down alarm generation and transfer is supported for all Gigabit Ethernet Interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit).

To configure this option, include the asynchronous-notification statement at the [edit interfaces ge- fpc/pic/port gigether-options] hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
    gigether-options {
        asynchronous-notification;
    }
}
```

SEE ALSO

asynchronous-notification

Configuring Link Down Notification for Optics Options Alarm or Warning

To configure this option, include the alarm or warning statement at the [edit interfaces ge- fpc/pic/port optics-options] hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
    optics-options {
```

```
alarm alarm-name {
        (syslog | link-down);
    }
    warning warning-name {
        (syslog | link-down);
     }
}
```

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
14.1R2	Starting with Junos OS Release 14.1R2, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC.
12.2	Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
12.2	Starting with Junos OS Release 12.2, WAN PHY mode is supported on the T4000 routers with the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP).
12.1R2	Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP).
12.1R2	Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the T4000 routers with the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP).

RELATED DOCUMENTATION

Configuring 100-Gigabit Ethernet MICs/PICs | 203

Configuring Gigabit Ethernet Policers | 255

Gigabit Ethernet Autonegotiation | 271

Configuring 40-Gigabit Ethernet PICs

IN THIS SECTION

- 40-Gigabit Ethernet PIC Overview | 200
- Configuring 40-Gigabit Ethernet PICs | 202

You can learn about the 40-Gigabit Ethernet PICs in this topic.

40-Gigabit Ethernet PIC Overview

The 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP) is a 1-port 40-Gigabit Ethernet Type 4 PIC with C form-factor pluggable transceiver (CFP) optics supported on T640, T1600, and T4000 routers. The 40-Gigabit Ethernet PIC occupies FPC slot 0 or 1 in the Type 4 FPC and it is similar to any regular PIC such as the 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP) PIC. The CFP information appears under the PIC information in the show command output.

The 40-Gigabit Ethernet PIC with CFP supports flexible Ethernet services encapsulation and MAC accounting.

MAC learning, MAC policing, and Layer 2 rewrite features are not supported.

The 40-Gigabit Ethernet PIC with CFP supports the following features:

- Encapsulation protocols such as:
 - Layer 2 protocols
 - Ethernet CCC, Ethernet TCC, and Ethernet VPLS
 - VLAN CCC
 - Extended VLAN TCC
 - VLAN VPLS
 - Flexible Ethernet service
 - Layer 3 protocols
 - IPv4
 - IPv6

- MPLS
- CFP Multi-Source Agreement (MSA)-compliant management data input/output (MDIO) control features (transceiver dependent).
- Graceful Routing Engine switchover (GRES) (in all PIC and chassis configurations).
- Interface creation:
 - When the PIC is brought online, the router creates one interface, et-x/y/0, where x represents the FPC slot number and y represents PIC slot number. The physical interface represents internal Ethernet Packet Forwarding Engines.
 - The FPC slot number ranges from 0 through 7 in T640, T1600, and T4000 routers. The PIC slot numbers are 0 and 1.
 - Packet Forwarding Engine 0 is the physical interface 0, and Packet Forwarding Engine 1 is the physical interface 1.
- 802.3 link aggregation:
 - The configuration of the 40-Gigabit Ethernet PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
 - An aggregate bundle that consists purely of 40-Gigabit Ethernet PICs supports a maximum of 40-Gigabit Ethernet links depending on the system implementation.

For Junos OS configuration information about this PIC, see "Configuring 40-Gigabit Ethernet PICs" on page 202. For hardware compatibility information, see the T1600 PICs Supported topic in the T1600 Core Router Hardware Guide hardware guide and the T640 PICs Supported topic in the T640 Core Router Hardware Guide hardware guide, and the T4000 PICs Supported topic in the T4000 Core Router Hardware Guide hardware guide.

SEE ALSO

T640 Core Router Hardware Guide	
T1600 Core Router Hardware Guide	
T4000 Core Router Hardware Guide	
TX Matrix Plus Router Hardware Guide	
T640 PICs Supported	
T1600 PICs Supported	
T4000 PICs Supported	

Configuring 40-Gigabit Ethernet PICs

You can configure the following features on the 40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP):

- Flexible Ethernet services encapsulation
- Source address MAC filtering
- Destination address MAC filtering
- MAC accounting for receive (Rx) and transmit (Tx)
- Multiple tag protocol ID (TPID) support
- Channels defined by two stacked VLAN tags
- Channels defined by flex-vlan-tagging
- IP service for stacked VLAN tags
- IP service for nonstandard TPID

The following features are not supported on the 40-Gigabit Ethernet PIC with CFP:

- MAC learning
- MAC policing
- Layer 2 rewrite

NOTE: Each 40-Gigabit Ethernet PIC with CFP creates a single et- physical interface in the Routing Engine and Packet Forwarding Engine.

The 40-Gigabit Ethernet PIC with CFP supports aggregated Ethernet configuration to achieve higher throughput capability, whereby the configuration is similar to the 1-Gigabit or 10-Gigabit aggregated Ethernet interface configuration. A maximum of 40-Gigabit Ethernet PIC links can be bundled into a single aggregated Ethernet configuration depending on the system implementation.

To configure the 40-Gigabit Ethernet PIC with CFP:

1. Perform the media configuration.

The command used to configure the media for the 40-Gigabit Ethernet PIC with CFP is the same as that for other Ethernet PICs, such as the 4-port 10-Gigabit Ethernet PIC.

2. Specify the logical interfaces.

A single physical interface is created when the 40-Gigabit Ethernet PIC with CFP is brought online (et-x/y/0), where x represents the FPC slot number and y represents the PIC slot number). For more

information, see Configuring Access Mode on a Logical Interface and Configuring a Logical Interface for Trunk Mode.

- 3. Configure the 802.3 link aggregation.
 - You must explicitly configure an aggregated interface on the 40-Gigabit Ethernet PIC with CFP that includes the 40-Gigabit Ethernet interfaces. For more information, see Configuring an Aggregated Ethernet Interface.
 - The configuration of the 40-Gigabit Ethernet PIC with CFP complies with the configuration of the 1-Gigabit Ethernet PIC, 10-Gigabit Ethernet PIC, and the aggregated Ethernet interfaces. In each aggregated bundle, Junos OS supports a maximum of 40-Gigabit Ethernet links. For more information, see Configuring an Aggregated Ethernet Interface and "10-port 10-Gigabit Ethernet LAN/WAN PIC Overview" on page 159.
- **4.** Configure the Packet Forwarding Engine features.

The 40-Gigabit Ethernet PIC with CFP supports all classification, firewall filters, queuing model, and rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see "Configuring Gigabit Ethernet Policers" on page 258, "Configuring Gigabit Ethernet Policers" on page 258, and Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview.

SEE ALSO

Configuring Gigabit Ethernet Policers | 258

Configuring Gigabit Ethernet Policers | 258

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers | 255

Gigabit Ethernet Autonegotiation | 271

Configuring 100-Gigabit Ethernet MICs/PICs

IN THIS SECTION

- 100-Gigabit Ethernet Interfaces Overview | 204
- MPC3E MIC Overview | 207

- 100-Gigabit Ethernet Type 4 PIC with CFP Overview | 208
- Configuring 100-Gigabit Ethernet Type 4 PIC With CFP | 211
- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 215
- 100-Gigabit Ethernet Type 5 PIC with CFP Overview | 218
- 100-Gigabit Ethernet Interfaces Interoperability | 220
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | 223
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 224
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | 226
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 227

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

100-Gigabit Ethernet Interfaces Overview

IN THIS SECTION

- MX Series 100-Gigabit Ethernet Interfaces | 204
- PTX Series 100-Gigabit Ethernet Interfaces | 205
- T Series 100-Gigabit Ethernet Interfaces | 206

MX Series 100-Gigabit Ethernet Interfaces

Table 21 on page 205 lists the 100-Gigabit Ethernet interfaces supported by MX Series routers.

Table 21: MX Series 100-Gigabit Ethernet Interfaces

Interface Module	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet MIC with CFP	MIC3-3D-1X100GE-CFP	MX240 MX480 MX960 MX2010 MX2020	100-Gigabit Ethernet MIC with CFP "MPC3E MIC Overview" on page 207
100-Gigabit Ethernet MIC with CXP	MIC3-3D-1X100GE-CXP	MX240 MX480 MX960 MX2010 MX2020	100-Gigabit Ethernet MIC with CXP "MPC3E MIC Overview" on page 207
100-Gigabit Ethernet ports on the MPC4E	MPC4E-3D-2CGE-8XGE	MX240 MX480 MX960 MX2010 MX2020	2x100GE + 8x10GE MPC4E
100-Gigabit Ethernet MIC with CFP2	MIC6-100G-CFP2	MX2010 MX2020	100-Gigabit Ethernet MIC with CFP2
100-Gigabit Ethernet MIC with CXP (4 Ports)	MIC6-100G-CXP	MX2010 MX2020	100-Gigabit Ethernet MIC with CXP (4 Ports)

PTX Series 100-Gigabit Ethernet Interfaces

Table 22 on page 206 lists the 100-Gigabit Ethernet interfaces supported by PTX Series routers.

Table 22: PTX Series 100-Gigabit Ethernet Interfaces

PIC	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet PIC with CFP	P1-PTX-2-100GE-CFP	PTX5000	100-Gigabit Ethernet PIC with CFP (PTX Series)
100-Gigabit Ethernet PIC with CFP2	P2-100GE-CFP2	PTX5000	100-Gigabit Ethernet PIC with CFP2 (PTX Series)
100-Gigabit Ethernet OTN PIC	P2-100GE-OTN	PTX5000	100-Gigabit Ethernet OTN PIC with CFP2 (PTX Series)
100-Gigabit DWDM OTN PIC	P1-PTX-2-100G-WDM	PTX5000 PTX3000	100-Gigabit DWDM OTN PIC (PTX Series)

T Series 100-Gigabit Ethernet Interfaces

Table 23 on page 206 lists the 100-Gigabit Ethernet interfaces supported by T Series routers.

Table 23: T Series 100-Gigabit Ethernet Interfaces

PIC	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet PIC with CFP (Type 4)	PD-1CE-CFP-FPC4	T1600 T4000	100-Gigabit Ethernet PIC with CFP (T1600 Router) 100-Gigabit Ethernet PIC with CFP (T4000 Router) "100-Gigabit Ethernet Type 4 PIC with CFP Overview" on page 208 "Configuring 100-Gigabit Ethernet Type 4 PIC With CFP" on page 211

Table 23: T Series 100-Gigabit Ethernet Interfaces (Continued)

PIC	Model Number	Routers Supported	For More Information
100-Gigabit Ethernet PIC with CFP (Type 5)	PF-1CGE-CFP	T4000	100-Gigabit Ethernet PIC with CFP (T4000 Router) "100-Gigabit Ethernet Type 5 PIC with CFP Overview" on page 218

SEE ALSO

MICs Supported by MX Series Routers

MPCs Supported by MX Series Routers

PICs Supported on the PTX Series

T1600 PICs Supported

T4000 PICs Supported

MPC3E MIC Overview

The MPC3E supports two separate slots for MICs. MICs provide the physical interface and are installed into the MPCs.

The MPC3E supports these MICs as field replaceable units (FRUs):

- 100-Gigabit Ethernet MIC with CFP (model number MIC3-3D-1X100GE-CFP)
- 100-Gigabit Ethernet MIC with CXP (model number MIC3-3D-1X100GE-CXP)
- 10-port 10-Gigabit Ethernet MIC with SFPP (model number MIC3-3D-10XGE-SFPP)
- 2-port 40-Gigabit Ethernet MIC with QSFP+ (model number MIC3-3D-2X40GE-QSFPP)

The MPC3E has two separate configurable MIC slots. Each MIC corresponds to a single PIC and the mapping between the MIC and PIC is 1 to 1 (one MIC is treated as one PIC). The MIC plugged into slot 0 corresponds to PIC 0 and the MIC plugged into slot 1 corresponds to PIC 2.

The MPC3E also supports these legacy MICs:

- 20-port Gigabit Ethernet MIC with SFP (model number MIC-3D-20GE-SFP)
- 2-port 10-Gigabit Ethernet MICs with XFP (model number MIC-3D-2XGE-XFP)

The 100-Gigabit Ethernet CFP MIC supports the IEEE standards—compliant 100BASE-LR4 interface, using the 100G CFP optical transceiver modules for connectivity. The 100-Gigabit Ethernet CXP MIC supports the 100BASE-SR10 interface, using 100-Gigabit CXP optical transceiver modules for connectivity. The 2-port 40-Gigabit Ethernet QSFPP MIC supports the 40BASE-SR4 interface and uses quad small form-factor pluggable (QSFPP) optical transceivers for connectivity. The 10-port 10-Gigabit Ethernet SFPP MIC uses SFP+ optical transceiver modules for connectivity.

For detailed information about each MIC, see 100-Gigabit Ethernet MIC with CFP, 100-Gigabit Ethernet MIC with CXP, 40-Gigabit Ethernet MIC with QSFP+. For information about supported hardware and transceivers, see MPC3E.

The MPC3E supports these features:

- Optical diagnostics and related alarms
- Virtual Router Redundancy Protocol (VRRP) support
- IEEE 802.1Q virtual LANs (VLANs) support
- Synchronous Ethernet
- Remote monitoring (RMON) and Ethernet statistics (EtherStats)
- Source MAC learning
- MAC accounting and policing—Dynamic local address learning of source MAC addresses
- Flexible Ethernet encapsulation
- Multiple Tag Protocol Identifiers (TPIDs)

NOTE: The MPC3E supports Ethernet interfaces only. SONET interfaces are not supported.

For information about the supported and unsupported Junos OS features for this MPC, see "Protocols and Applications Supported by the MPC3E (MX-MPC3E)" in the MX Series Interface Module Reference.

SEE ALSO

Protocols and Applications Supported by the MPC3E on MX Series Routers

MX Series Interface Module Reference

100-Gigabit Ethernet Type 4 PIC with CFP Overview

The 100-Gigabit Ethernet PIC (model number PD-1CE-CFP-FPC4) is a 1-port 100-Gigabit Ethernet Type 4 PIC with 100-gigabit small form-factor pluggable (CFP) transceiver. This PIC is available only as packaged in an assembly with the T1600-FPC4-ES FPC. The 100-Gigabit Ethernet PIC occupies PIC slots 0 and 1 in the T1600-FPC4-ES FPC. For information about supported transceivers and hardware, see 100-Gigabit Ethernet PIC with CFP (T1600 Router).

The 100-Gigabit Ethernet PIC supports flexible encapsulation and MAC accounting.

MAC learning, MAC policing, and Layer 2 rewrite functionality are not supported.

The ingress flow can be filtered based on the VLAN source and destination addresses. Ingress frames can also be classified according to VLAN, stacked VLAN, source address, VLAN source address, and stacked VLAN source address. VLAN manipulation on egress frames are supported on both outer and inner VLAN tags.

The following features are supported:

- The following encapsulation protocols are supported:
 - Layer 2 protocols
 - Ethernet CCC, Ethernet TCC, Ethernet VPLS
 - VLAN CCC
 - Extended VLAN TCC
 - VLAN VPLS
 - Flexible Ethernet service
 - Layer 3 protocols
 - IPv4
 - Ipv6
 - MPLS
- CFP MSA compliant MDIO control features (transceiver dependent).
- Graceful Routing Engine switchover (GRES) is supported in all PIC and chassis configurations.
- Interface creation:
 - When the PIC, is brought online, the router creates two 50 gigabit capable interfaces, et-x/0/0:0 and et-x/0/0:1, where x represents the FPC slot number. Each physical interface represents two internal 50 gigabit Ethernet Packet Forwarding Engines. Two logical interfaces are configured under each physical interface.

- Packet Forwarding Engine 0 is physical interface 0, Packet Forwarding Engine 1 is physical interface 1
- 802.3 link aggregation:

Same rate or same mode link aggregation:

- Two logical interfaces are created for each 100-Gigabit Ethernet PIC. To utilize bandwidth beyond 50 gigabits per second, an aggregate interface must be explicitly configured on the 100-Gigabit Ethernet PIC that includes the two 50 gigabit interfaces.
- Each 100 gigabit Ethernet aggregate consumes one of the router-wide aggregated Ethernet device pools. The number of 100-Gigabit Ethernet PICs cannot exceed the router-wide limit, which is 128 for Ethernet.
- In each aggregate bundle, each 100-Gigabit Ethernet PIC consumes two members. Hence, an aggregate bundle that consists purely of 100-Gigabit Ethernet PICs supports a maximum of half of the software limit for the number of members. Therefore, with a maximum of 16 links, up to 8 100-Gigabit Ethernet links are supported.
- Combining 100-Gigabit Ethernet PICs into aggregate interfaces with other Ethernet PICs is not permitted. However, other Ethernet PICs can also be configured within the same T1600 with 100-Gigabit Ethernet PICs, and used in separate aggregate interfaces.
- Multiple (Juniper Networks) Type 4 100-Gigabit Ethernet PICs on a T1600 router can be
 combined into a static aggregated Ethernet bundle to connect to a different type of 100 gigabit
 Ethernet PIC on a remote router (Juniper Networks or other vendors). LACP is not supported in
 this configuration.

Mixed rate or mixed mode link aggregation:

- Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC.
- Static link protection and Link Aggregation Control Protocol (LACP) is supported on mixed aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC.
- When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC, ensure that
 you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC to the
 aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be
 included in the same aggregated Ethernet bundle.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit
 Ethernet PIC, the packet loss performance value is twice the original value because of the two 50Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC.

- Software Packet Forwarding Engine—Supports all Gigabit Ethernet PIC classification, *firewall filter*, queuing model, and rewrite functionality.
- Egress traffic performance—Maximum egress throughput is 100 gigabits per second on the physical interface, with 50 gigabits per second on the two assigned logical interfaces.
- Ingress traffic performance—Maximum ingress throughput is 100 gigabits per second on the physical interface, with 50 gigabits per second on the two assigned logical interfaces. To achieve 100 gigabits per second ingress traffic performance, use one of the interoperability modes described below. For example, if VLAN steering mode is not used when connecting to a remote 100 gigabits per second interface (that is on a different 100 gigabits per second PIC on a Juniper Networks router or a different vendor's equipment), then all ingress traffic will try to use one of the 50 gigabits per second Packet Forwarding Engines, rather than be distributed among the two 50 gigabits per second Packet Forwarding Engines, resulting in a total of 50 gigabits per second ingress performance.
- Interoperability modes—The 100-Gigabit Ethernet PIC supports interoperability with through configuration in one of the following two forwarding option modes:
 - *SA multicast mode*—In this mode, the 100-Gigabit Ethernet PIC supports interconnection with other Juniper Networks 100-Gigabit Ethernet PICs (Model: PD-1CE-CFP) interfaces only.
 - VLAN steering mode—In this mode, the 100-Gigabit Ethernet Type 4 PIC with CFP supports interoperability with 100 gigabit Ethernet interfaces from other vendors only.

SEE ALSO

T1600 Core Router Hardware Guide

100-Gigabit Ethernet PIC with CFP (T1600 Router)

100-Gigabit Ethernet PIC with CFP (T4000 Router)

Configuring 100-Gigabit Ethernet Type 4 PIC With CFP

You can configure the following features on the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4):

- Flexible Ethernet services encapsulation
- Source address MAC filtering
- Destination address MAC filtering
- MAC accounting in RX
- Channels defined by two stacked VLAN tags

- · Channels defined by flex-vlan-tagging
- IP service for stacked VLAN tags
- Layer 2 rewrite

The following features are not supported on the 100-Gigabit Ethernet Type 4 PIC with CFP:

- Multiple TPID
- IP service for non-standard TPID
- MAC learning
- MAC policing

NOTE:

- For the 100-Gigabit Ethernet Type 4 PIC with CFP, only the PICO online and offline CLI commands are supported. The PIC1 online and offline CLI commands are not supported.
- Each 100-Gigabit Ethernet Type 4 PIC with CFP creates two et- physical interfaces, defined as 50-gigabit physical interfaces in the Routing Engine and Packet Forwarding Engine. By default, these are independent physical interfaces and are not configured as an aggregated Ethernet interface.

To configure a 100-Gigabit Ethernet Type 4 PIC with CFP:

1. Perform the media configuration:

The 100-Gigabit Ethernet Type 4 PIC with CFP features a 100 gigabit per second pipe. The media-related configuration commands for $et-x/\theta/\theta$: 0 and $et-x/\theta/\theta$:1 must both be configured at the same time and configured with the same value, otherwise the commit operation fails.

When configuring to activate or deactivate the interface, if the interface contains the described media-related configuration, it must activate and deactivate both units 0 and 1 at the same time, otherwise the commit operation fails.

The following media configuration commands have the above described restriction:

- # set interfaces et-x/0/0:1 disable
- # set interfaces et-x/0/0:1 gigether-options loopback
- # set interfaces et-x/0/0:1 mtu yyy

Due to an MTU restriction, the vlan-tagging and flexible-vlan-tagging configuration on et-x/0/0:0 and et-x/0/0:1 must be same, otherwise the commit operation fails.

2. Specify the logical interfaces:

- **a.** Two physical interfaces are created when the 100-Gigabit Ethernet Type 4 PIC with CFP is brought online (et-x/0/0:0 and et-x/0/0:1, where x represents the FPC slot number). Each physical interface represents two internal 50-gigabit Ethernet Packet Forwarding Engines.
- **b.** Two logical interfaces are configured under each physical interface: Packet Forwarding Engine 0 is physical interface 0 and Packet Forwarding Engine 1 is physical interface 1.
- **3.** Configure the 802.3 link aggregation:
 - a. The 100-Gigabit Ethernet PIC supports aggregated Ethernet configuration to achieve higher throughput capability, whereby configuration is similar to the 1G/10G aggregated Ethernet interface configuration.
 - **b.** Two physical interfaces are created for each 100-Gigabit Ethernet Type 4 PIC with CFP. To utilize bandwidth beyond 50 gigabits, a same rate and same mode aggregated Ethernet interface must be explicitly configured on the 100-Gigabit Ethernet Type 4 PIC with CFP that includes these two 50-gigabit interfaces.
 - c. Each 100-Gigabit Ethernet Type 4 PIC with CFP aggregate consumes one of the router-wide aggregated Ethernet device pools. In Junos OS with 100-Gigabit Ethernet PICs, you cannot exceed the router limit of 128 Ethernet PICs.
 - d. In each aggregated bundle, each 100-Gigabit Ethernet Type 4 PIC with CFP consumes two aggregate members. Hence, an aggregated bundle consisting of only one 100-Gigabit Ethernet Type 4 PIC with CFP supports only up to half of the Junos OS limit for the number of members. The Junos OS supports a maximum of 16 links for up to 8 100-Gigabit Ethernet Type 4 PIC with CFP links.

NOTE: The 100-Gigabit Ethernet Type 4 PIC with CFP has the following restrictions for same rate and same mode aggregated Ethernet configuration:

- Both physical interfaces belonging to the same 100-Gigabit Ethernet PIC must be included in the same aggregated Ethernet physical interfaces. The aggregation of the 100-Gigabit Ethernet PIC interface is always an even number of physical interfaces.
- The 100-Gigabit Ethernet PIC physical interface cannot be configured in the aggregated interface with any other type of physical interface.
- The maximum supported number of aggregated 100-Gigabit Ethernet PIC interfaces is half of the number that the Junos OS supports for 1G/10G aggregated Ethernet. For example, if Junos OS supports 16 ports of 10-gigabit Ethernet aggregation, it supports 8 ports of 100-Gigabit Ethernet PIC aggregation. This is because each port of the 100-

Gigabit Ethernet PIC port using 2 physical interfaces (et-x/0/0:0 and et-x/0/0:1), where each physical interface represents 50 gigabits of traffic capacity.

e. Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC. When configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC, ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.

NOTE: The 100-Gigabit Ethernet Type 4 PIC with CFP has the following restrictions for mixed rate and mixed mode aggregated Ethernet configuration:

- A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
- Traffic distribution is based on the hash calculated on the egress packet header. Hash
 range is fairly distributed according to member links' speed. This guarantees hash fairness
 but it does not guarantee fair traffic distribution depending on the rate of the egress
 streams.
- Packets are dropped when the total throughput of the hash flow exiting a member link (or
 multiple hash flows exiting a single member link) exceeds the link speed of the member
 link. This can happen when egress member link changes because of a link failure and the
 hash flow switches to a member link of speed that is less than the total throughput of the
 hash flow.
- Rate-based CoS components such as scheduler, shaper, and policer are not supported on mixed rate aggregated Ethernet links. However, the default CoS settings are supported by default on the mixed rate aggregated Ethernet links.
- Load balancing is performed at the ingress Packet Forwarding Engine. Therefore, you must ensure that the egress traffic on the aggregated Ethernet link enters through the hardware platforms that support mixed aggregated Ethernet bundles.
- Mixed aggregated Ethernet links can interoperate with non-Juniper Networks aggregated Ethernet member links provided that mixed aggregated Ethernet load balancing is configured at egress.
- 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.

- Changing the edit interfaces aex aggregated-ether-options link-speed configuration of a mixed aggregated Ethernet link, which is configured on the supported interfaces of on T640, T1600, T4000, and TX Matrix Plus routers, leads to aggregated Ethernet link flapping.
- When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC, changing aggregated Ethernet link protection configurations leads to aggregated Ethernet link flapping.
- For a single physical link event of an aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC, the packet loss performance value is twice the original value because of the *two* 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP.
- The show interfaces aex command displays the link speed of the aggregated Ethernet interface, which is the sum of the link speeds of all the active member links.
- **4.** Configure the Packet Forwarding Engine features:
 - a. The 100-Gigabit Ethernet Type 4 PIC with CFP supports all classification, firewall filters, queuing model, and rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see "Configuring Gigabit Ethernet Policers" on page 258, "Configuring Gigabit Ethernet Policers" on page 258, and Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview.

NOTE: When using the show interfaces extensive command with a 100-Gigabit Ethernet Type 4 PIC with CFP, the "Filter statistics" section will not be displayed because the hardware does not include those counters.

SEE ALSO

Configuring Gigabit Ethernet Policers | 258

Configuring Gigabit Ethernet Policers | 258

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP

In Junos OS Release 10.4 and later, you can configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with routers using 100 gigabit Ethernet interfaces from other vendors by using the forwarding-mode statement with the vlan-steering option at the [edit chassis fpc *slot* pic *slot*] hierarchy level. On ingress, the router compares the outer VLAN ID against the user-defined VLAN ID and VLAN mask combination and steers the packet accordingly. You can program a custom VLAN ID and corresponding mask for PFE0.

General information on the VLAN steering mode:

- In VLAN steering mode, the SA multicast parameters are not used for packet steering.
- In SA multicast bit steering mode, the VLAN ID and VLAN masks are not used for packet steering.
- Configuration to set the packet distribution mode and VLAN steering rule is done through CLI commands. Both CLI commands result in a PIC reboot.
- There are three possible tag types of ingress packet:
 - Untagged ingress packet—The packet is sent to PFE1.
 - Ingress packet with one VLAN—The packet is forwarded to the corresponding PFE based on the VLAN ID.
 - Ingress packet with two VLANs—The packet is forwarded to the corresponding PFE based on the outer VLAN ID.
- If no VLAN rule is configured, all tagged packets are distributed to PFEO.
- VLAN rules describe how the router distributes packets. Two VLAN rules are provided by the CLI:
 - Odd-Even rule—Odd number VLAN IDs go to PFE1; even number of VLAN IDs go to PFE0.
 - Hi-Low rule—VLAN IDs 1 through 2047 go to PFE0; VLAN IDs 2048 through 4096 go to PFE1.
- When the 100-Gigabit Ethernet Type 4 PIC with CFP is configured in VLAN steering mode, it can be configured in a two physical interfaces mode or in aggregate Ethernet (AE) mode:
 - Two physical interfaces mode—When the PIC is in the two physical interfaces mode, it creates the physical interfaces et-x/0/0:0 and et-x/0/0:1. Each physical interface can configure its own logical interface and VLAN. The CLI enforces the following restrictions at the commit time:
 - The VLAN ID configuration must comply with the selected VLAN rule.
 - The previous restriction implies that the same VLAN ID cannot be configured on both physical interfaces.
 - AE mode—When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one AE physical interface. The PIC egress traffic is based on an AE internal hash algorithm. The PIC ingress traffic steering is based on the customized VLAN ID rule. The CLI enforces the following restrictions at the commit time:
 - The PICs AE working in VLAN steering mode includes both links of that PIC, and only the links
 of that PIC.
 - The PIC AE working in SA multicast steering mode can include more than one 100-Gigabit Ethernet Type 4 PIC with CFP to achieve more than 100 gigabit Ethernet capacity.

To configure SA multicast mode, use the set chassis fpc *slot* pic *slot* forwarding-mode sa-multicast command.

SA Multicast Mode

To configure SA multicast mode on a Juniper Networks 100-Gigabit Ethernet Type 4 PIC with CFP in FPC 0, PIC 0 for interconnection with another Juniper Networks 100-Gigabit Ethernet PIC, use the set chassis fpc *slot* pic *slot* forwarding-mode sa-multicast command. You can use the show forwarding-mode command to view the resulting configuration, as follows:

```
[edit chassis fpc slot pic slot]
user@host# show forwarding-mode
forwarding-mode {
    sa-multicast;
}
```

VLAN Steering Mode

To configure the Juniper Networks 100-Gigabit Ethernet Type 4 PIC with CFP for VLAN steering mode for interoperation with a 100 gigabit Ethernet interface from another vendor's router, use the set chassis fpc *slot* pic *slot* forwarding-mode vlan-steering command with the vlan-rule (high-low | odd-even) statement. You can use the show forwarding-mode command to view the resulting configuration, as follows:

```
[edit chassis fpc slot pic slot]
user@host# show forwarding-mode
forwarding-mode {
    vlan-steering {
        vlan-rule odd-even;
    }
}
```

SEE ALSO

```
forwarding-mode (100-Gigabit Ethernet)
sa-multicast (100-Gigabit Ethernet)
vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP)
```

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP)

100-Gigabit Ethernet Type 5 PIC with CFP Overview

The 100-Gigabit Ethernet PIC is a 1-port 100-Gigabit Ethernet Type 5 PIC with C form-factor pluggable transceiver (CFP) with model number PF-1CGE-CFP.

The following features are supported on 100-Gigabit Ethernet Type 5 PIC with CFP:

- Access to all 100-Gigabit Ethernet port counters through SNMP.
- Logical interface-level MAC filtering, accounting, policing, and learning for source media access control (MAC).
- Channels defined by two stacked VLAN tags.
- Channels defined by flex-vlan-tagging.
- IP service for stacked VLAN tags.
- Defining the rewrite operation to be applied to the incoming and outgoing frames on logical interfaces on this PIC.

NOTE: Only the Tag Protocol Identifier (TPID) 0x8100 is supported.

- Interface encapsulations, such as the following:
 - untagged—Default encapsulation, when other encapsulation is not configured.
 - You can configure only one logical interface (unit 0) on the port.
 - You cannot include the vlan-id statement in the configuration of the logical interface.
 - vlan-tagging—Enable VLAN tagging for all logical interfaces on the physical interface.
 - stacked-vlan-tagging—Enable stacked VLAN tagging for all logical interfaces on the physical interface.
 - ethernet-ccc—Ethernet cross-connect.
 - ethernet-tcc—Ethernet translational cross-connect.
 - vlan-ccc—802.1Q tagging for a cross-connect.
 - vlan-tcc—Virtual LAN (VLAN) translational cross-connect.
 - extended-vlan-ccc—Standard TPID tagging for an Ethernet cross-connect.
 - extended-vlan-tcc—Standard TPID tagging for an Ethernet translational cross-connect.

- flexible-ethernet-services—Allows per-unit Ethernet encapsulation configuration.
- ethernet-vpls—Ethernet virtual private LAN service.
- vlan-vpls—VLAN virtual private LAN service.
- The following Layer 3 protocols are also supported:
 - IPv4
 - IPv6
 - MPLS
- CFP Multi-Source Agreement (MSA) compliant Management Data Input/Output (MDIO) control features (transceiver dependent).
- 802.3 link aggregation:
 - The configuration of the 100-Gigabit Ethernet Type 5 PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
- Interoperability mode—Interoperability with the 100-Gigabit Ethernet Type 4 PIC with CFP through configuration in sa-multicast forwarding mode.
- Juniper Networks enterprise-specific Ethernet Media Access Control (MAC) MIB
- The 100-Gigabit Ethernet Type 5 PIC with CFP supports all Gigabit Ethernet PIC classification, firewall filters, queuing model, and Layer 2 rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see "Configuring Gigabit Ethernet Policers" on page 258, "Configuring Gigabit Ethernet Policers" on page 258, and Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview.
- A Type 5 FPC can support up to two 100-Gigabit Ethernet PICs. Both the PICs (that is, PIC 0 and PIC
 1) can be offline or online independently.

The following features are not supported on the 100-Gigabit Ethernet Type 5 PIC with CFP:

MAC filtering, accounting, and policing for destination MAC at the logical interface level.

NOTE: Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This configuration enables the OSPF protocol to work.

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.

- Multiple TPIDs.
- IP service for nonstandard TPID.

Table 24 on page 220 lists the capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP.

Table 24: Capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP

Capability	Support
Maximum logical interfaces per PIC	4093
Maximum logical interfaces per port	For IPv4 the limit is 4093. For IPv6 the limit is 1022.

SEE ALSO

Configuring Gigabit Ethernet Policers | 258

Configuring Gigabit Ethernet Policers | 258

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

100-Gigabit Ethernet Interfaces Interoperability

IN THIS SECTION

- Interoperability of the MIC-3D-1X100GE-CFP MIC with PICs on Other Routers | 221
- Interoperability of the MPC4E-3D-2CGE-8XGE MPC with PICs on Other Routers | 221
- Interoperability of the P1-PTX-2-100GE-CFP PIC with PICs on Other Routers | 221
- Interoperability of the PD-1CE-CFP-FPC4 PIC with PICs or MICs on Other Routers | 222

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.

Interoperability of the MIC-3D-1X100GE-CFP MIC with PICs on Other Routers

Table 25 on page 221 lists the Interoperability with the 100-Gigabit Ethernet MIC with CFP.

Table 25: 100-Gigabit Ethernet MIC with CFP (MIC3-3D-1X100GE-CFP) Interoperability

Interoperates with		For More Information
T Series	100-Gigabit Ethernet PIC with CFP (Type 4) (PD- 1CE-CFP-FPC4)	Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode

Interoperability of the MPC4E-3D-2CGE-8XGE MPC with PICs on Other Routers

Table 26 on page 221 lists the Interoperability with the MPC4E.

Table 26: MPC4E Interoperability

Interoperates with		For More Information
T Series	100-Gigabit Ethernet PIC with CFP (Type 4) (PD-1CECFP- FPC4)	Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode

Interoperability of the P1-PTX-2-100GE-CFP PIC with PICs on Other Routers

Table 27 on page 221 lists the Interoperability with 100-Gigabit Ethernet PIC with CFP (Type 5).

Table 27: 100-Gigabit Ethernet PIC with CFP (Type 5) (P1-PTX-2-100GE-CFP) Interoperability

Interoperates with	For More Information

T Series		"Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP" on page 226
	CFP-FPC4)	"Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4" on page 227

Interoperability of the PD-1CE-CFP-FPC4 PIC with PICs or MICs on Other Routers

Table 28 on page 222 lists the 100-Gigabit Ethernet PIC with CFP (Type 4).

Table 28: 100-Gigabit Ethernet PIC with CFP (Type 4) PD-1CE-CFP-FPC4 Interoperability

Interoperates with		For More Information
T Series	100-Gigabit Ethernet PIC with CFP (Type 5) (PF- 1CGE-CFP)	"Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4" on page 224 forwarding-mode sa-multicast
MX Series	100-Gigabit Ethernet MIC with CFP (MIC3-3D- 1X100GE- CFP)	Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode
	100-Gigabit Ethernet ports on the MPC4E	Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode
PTX Series	100-Gigabit Ethernet PIC with CFP (Type 5) (P1- PTX-2-100GE-CFP)	"Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP" on page 226 "Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4" on page 227

SEE ALSO

Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP by:

- Enabling source address (SA) multicast bit steering mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP.
- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.

SA multicast mode uses the multicast bit in the source MAC address for packet steering. By default, the SA multicast bit is set to 0 for all packets sent by the 100-Gigabit Ethernet PIC PF-1CGE-CFP. The 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 looks at the bit and forwards the packets to either Packet Forwarding Engine 0 or Packet Forwarding Engine 1. When the PIC sends out a packet, the multicast bit is set based on the egress Packet Forwarding Engine number (0 or 1).

The default packet steering mode for PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.

PD-1CE-CFP-FPC4 uses two 50 Gpbs Packet Forwarding Engines to achieve 100 Gbps throughput. The 50-Gigabit Ethernet physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both the physical interfaces. You must configure the physical interfaces on PD-1CE-CFP-FPC4 in static link aggregation group (LAG) mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to the 100-Gigabit Ethernet PIC PF-1CGE-CFP instead of two independent 50-Gigabit Ethernet interfaces.

NOTE: If you try to enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP without configuring PD-1CE-CFP-FPC4 (with two 50-Gigabit Ethernet interfaces) in static LAG mode, then there are issues in forwarding or routing protocols. For example, if you create two untagged logical interfaces—one each on the two 50-Gigabit Ethernet interfaces—on PD-1CE-CFP-FPC4 and one untagged *logical interface* on PF-1CGE-CFP, then PF-1CGE-CFP does not learn about one of the 50-Gigabit Ethernet interfaces on PD-1CE-CFP-FPC4.

SEE ALSO

forwarding-mode

sa-multicast

100-Gigabit Ethernet PIC with CFP (T1600 Router)

100-Gigabit Ethernet PIC with CFP (T4000 Router)

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4

IN THIS SECTION

Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP | 224

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP by performing the following tasks:

Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP, you need to enable source address (SA) multicast bit steering mode on PF-1CGE-CFP.

To configure SA multicast mode on PF-1CGE-CFP:

1. Specify the FPC and PIC information on the chassis.

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

For example:

```
[edit ]
user@host# edit chassis fpc 1 pic 0
```

2. Configure the interoperation mode (SA multicast bit steering mode).

```
[edit chassis fpc slot pic slot]
user@host# set forwarding-mode sa-multicast
```

For example:

```
[edit fpc 1 pic 0]
user@host# set forwarding-mode sa-multicast
```

3. Verify the configuration.

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

fpc 1 {
    pic 0 {
        forwarding-mode {
            sa-multicast;
        }
    }
}
```

NOTE: The default packet steering mode for the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 is SA multicast bit mode. No SA multicast configuration is required to enable this mode.

SEE ALSO

100-Gigabit Ethernet PIC with CFP (T1600 Router)

100-Gigabit Ethernet PIC with CFP (T4000 Router)

RELATED DOCUMENTATION

100-Gigabit Ethernet PIC with CFP (T1600 Router)

100-Gigabit Ethernet PIC with CFP (T4000 Router)

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP

You can enable interoperability between the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 and the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP by:

- Configuring the two 50-Gigabit Ethernet physical interfaces on the 100-Gigabit Ethernet PIC
 PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface.
- Configuring source address (SA) multicast bit steering mode on the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP.

SA multicast bit steering mode uses the multicast bit in the source MAC address for packet steering.

NOTE: When SA multicast bit steering mode is configured on a PTX Series Packet Transport Router 100-Gigabit Ethernet port, VLANs are not supported for that port.

The 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 uses two 50-Gpbs Packet Forwarding Engines to achieve 100-Gbps throughput. The 50-Gigabit Ethernet physical interfaces are created when the 100-Gigabit Ethernet PIC is plugged in. The two physical interfaces are visible and configuration is allowed on both the physical interfaces. You must configure the physical interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 in static link aggregation group (LAG) mode without enabling Link Aggregation Control Protocol (LACP). This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP.

On the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4, ingress packets are forwarded to either Packet Forwarding Engine number 0 or 1 based on the SA multicast bit in the received packet. The SA multicast bit of egress packets is set based on whether the packet is forwarded from Packet Forwarding Engine number 0 or 1. As the default packet steering mode is SA multicast bit steering mode, no configuration is necessary to enable this mode.

On the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP, the SA multicast bit is ignored in ingress packets. When SA multicast bit steering mode is enabled, the SA multicast bit in the egress packets is set to 0 or 1 based on the flow hash value that is computed internally by the Packet Forwarding Engine complex for each packet. No CLI configuration is required to generate the flow hash value as this computation is done automatically. The flow hash algorithm uses fields in the packet header to compute the flow hash value. By default, the SA multicast bit is set to 0 in egress packets. You must configure SA multicast bit steering mode to enable interoperability with the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4.

NOTE: If you try to enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP without configuring PD-1CE-CFP-FPC4 (with two 50-Gigabit Ethernet interfaces) in static LAG mode, then there are issues in forwarding or routing protocols. For example, if you create two untagged logical interfaces—one each on the two 50-Gigabit Ethernet interfaces—on the PD-1CE-CFP-FPC4 and one untagged *logical interface* on the P1-PTX-2-100GE-CFP, then P1-PTX-2-100GE-CFP does not learn about one of the 50-Gigabit Ethernet interfaces on PD-1CE-CFP-FPC4.

SEE ALSO

sa-multicast

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4

IN THIS SECTION

- Configuring SA Multicast Bit Steering Mode on 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP | 227
- Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface | 229

You can enable interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP by performing the following tasks:

Configuring SA Multicast Bit Steering Mode on 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP, you must enable source address (SA) multicast bit steering mode on P1-PTX-2-100GE-CFP.

NOTE: When you configure the SA multicast bit steering mode on the PTX Series PIC P1-PTX-2-100GE-CFP, we recommend that you do not configure the PIC ports as member links of an aggregated Ethernet interface because this prevents load balancing on the peering T Series

PIC PD-1CE-CFP-FPC4. This T Series PIC must be in aggregated Ethernet mode to share bandwidth between its two 50-Gigabit Ethernet interfaces.

To configure SA multicast bit steering mode on the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP:

1. Specify the FPC, PIC, and port information on the chassis.

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

For example:

```
[edit ]
user@host# edit chassis fpc 1 pic 0 port 0
```

2. Configure the interoperation mode (SA multicast bit steering mode).

```
[edit chassis fpc 1 pic 0]
user@host# set forwarding-mode sa-multicast
```

3. Verify the configuration.

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

fpc 1 {
    pic 0 {
        forwarding-mode {
            sa-multicast;
            }
        }
    }
}
```

NOTE: As the default packet steering mode for the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 is SA multicast bit steering mode, no configuration is necessary to enable this mode.

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface

To enable the interoperability between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP or P1-PTX-2-100GE-CFP, you need to configure the two 50-Gigabit Ethernet physical interfaces on PD-1CE-CFP-FPC4 as one aggregated Ethernet physical interface. This ensures that a single 100-Gigabit aggregated interface is visible on the link connecting to PF-1CGE-CFP or P1-PTX-2-100GE-CFP instead of two independent 50-Gigabit Ethernet interfaces.

When the PIC is in aggregated Ethernet mode, the two physical interfaces on the same PIC are aggregated into one aggregated Ethernet physical interface. When the PIC is configured with two physical interfaces, it creates the physical interfaces et-*fpc*/*pic*/0:0 and et-*fpc*/*pic*/0:1, where *fpc* is the FPC slot number and *pic* is the PIC slot number. For example, to configure two physical interfaces for PIC slot 0 in FPC slot 5:

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

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

For example:

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

2. Specify the members to be included within the aggregated Ethernet bundle.

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

The following example shows how to configure two physical interfaces for PIC 0 on a T1600 router.

```
[edit interfaces ]
user@host# set et-5/0/0:0 gigether-options 802.3ad ae0
user@host# set et-5/0/0:1 gigether-options 802.3ad ae0
```

3. Verify the configuration at the chassis.

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

```
aggregated-devices {
    ethernet {
        device-count 1;
    }
}
```

4. Verify the configuration at the interface.

```
[edit ]
user@host# show interfaces
et-5/0/0:0 {
    gigether-options {
       802.3ad ae0;
    }
}
et-5/0/0:1 {
    gigether-options {
       802.3ad ae0;
    }
}
```

SEE ALSO

Configuring Junos OS for Supporting Aggregated Devices

802.3ad

RELATED DOCUMENTATION

sa-multicast

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
13.2	Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC.

RELATED DOCUMENTATION

Configuring Gigabit Ethernet Policers | 255

Gigabit Ethernet Autonegotiation | 271

Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks

IN THIS SECTION

- Transporting Legacy Traffic over Packet Switched Networks | 231
- Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH Traffic | 234
- Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy SDH Traffic | 240

This topic describes how to transport legacy TDM traffic over Packet switched networks using Smart SFP transceivers.

Transporting Legacy Traffic over Packet Switched Networks

IN THIS SECTION

- Smart SFP Transceivers for Transporting PDH Traffic over PSNs Overview | 232
- Smart SFP Transceivers for Transporting SDH Traffic over PSNs Overview | 233
- Benefits of Smart SFP Transceivers | 234

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*. Examples of PDH interfaces are E1, T1, and DS3. In SDH traffic, the bit streams are of the same bit rate but are derived from a common clock and are thus synchronous. Examples of SDH interfaces are STM1, STM4, and STM16. Based on the type of legacy TDM traffic, PDH or SDH, you can choose Smart SFP optics to convert the legacy packets to Ethernet frames that can be transported over PSNs.

Smart SFP Transceivers for Transporting PDH Traffic over PSNs Overview

Junos OS supports the following three smart SFP transceivers on MX Series 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)

On your MX Series routers, 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. The MX104 routers with in-built 1G and 10G MPCs also 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 MX Series 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 MX104 routers with in-built 1G and 10G MPCs also support the smart SFP transceivers.

NOTE: 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.

Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH Traffic

IN THIS SECTION

- Requirements for Configuration of the Smart-SFPs on MX Series Routers | 234
- Overview | 235
- Configuring the DS3 Smart SFP | 235
- Verification | 237

Requirements for Configuration of the Smart-SFPs on MX Series Routers

This example uses the following hardware and software components:

- Junos OS Release 19.4R1 or later for MX Series routers
- A single MX480 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 an MX480 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.

Configuring the DS3 Smart SFP

IN THIS SECTION

• Procedure | 235

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 | 238
- Verifying the DS3 Smart SFP Defects on the Interface | 239

To verify that the DS3 Smart SFP is configured on the MX480 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 MX480 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
                                                 1187124
    Rx unicast frames
    Rx multicast frames
    Rx broadcast frames
                                                       0
    Rx fragments
                                                       0
    Rx undersize frames
                                                       0
    Rx oversize frames
    Rx invalid vlan mismatch frames
    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:		Counte	ers
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 Encapsul	ation: ENET2		
Input packets : 1384454023			
utput packets: 0			
Protocol multiservice, MTU: Unlimited			

Meaning

The DS3 Smart SFP is configured on the MX480 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 MX480 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
   Encapulation 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 MX480 router and you can view the DS3 Smart SFP defects.

Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy SDH Traffic

IN THIS SECTION

- Requirements for Configuration of the Smart-SFPs on MX Series Routers | 241
- Overview | 241
- Configuring the STM1 Smart SFP | 241
- Verification | 243

Requirements for Configuration of the Smart-SFPs on MX Series Routers

This example uses the following hardware and software components:

- Junos OS Release 19.4R1 or later for MX Series routers
- A single MX480 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 an MX480 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.

Configuring the STM1 Smart SFP

IN THIS SECTION

• Procedure | 241

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 | 244
- Verifying the STM1 Smart SFP Defects on the Interface | 245

To verify that the STM1 Smart SFP is configured on the MX480 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 MX480 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
   Rx unicast frames
                                                 1187124
   Rx multicast frames
   Rx broadcast frames
                                                       0
   Rx fragments
                                                       0
   Rx undersize frames
                                                       0
   Rx oversize frames
   Rx invalid vlan mismatch frames
   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:		Count	ers	
BiPolarVariations/Excessive zero error	S	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 Encapsu	lation: ENET2			
Input packets : 1384454023				
Output packets: 0				
Protocol multiservice, MTU: Unlimited				

Meaning

The STM1 Smart SFP is configured on the MX480 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
    Encapulation 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 MX480 router and you can view the STM1 Smart SFP defects.

Configuring Layer 2 Overhead Attribute in Interface Statistics

IN THIS SECTION

- Accounting of the Layer 2 Overhead Attribute in Interface Statistics | 247
- Configuring Layer 2 Overhead Accounting in Interface Statistics | 250
- Verifying the Accounting of Layer 2 Overhead in Interface Statistics | 252

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. Use this topic to understand more 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,

Accounting of the Layer 2 Overhead Attribute in Interface Statistics

IN THIS SECTION

Guidelines for Configuring the Computation of Layer 2 Overhead in Interface Statistics | 249

On MX Series and T Series routers, 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 MX Series routers. 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. 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).

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 T Series 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 29 on page 248 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 29: 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)
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

Table 29: Adjustment Bytes for Logical Interfaces over Ethernet Interfaces (Continued)

Encapsulation Type on Logical Interfaces	Number of Adjustment Bytes	Description
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 for Configuring 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.

SEE ALSO

account-layer2-overhead

Configuring Layer 2 Overhead Accounting in Interface Statistics

IN THIS SECTION

Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level | 251

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.

Enabling 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
```

SEE ALSO

account-layer2-overhead

Verifying the Accounting of Layer 2 Overhead in Interface Statistics

IN THIS SECTION

- Purpose | 252
- Action | 252

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:

NOTE: 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:
   Output packets:
                                362633
  Multicast statistics:
   IPV4 multicast statistics:
    Input bytes :
                                                            0 bps
                                 112048
    Output bytes :
                                20779920
                                                            0 bps
    Input packets:
                                    1801
                                                            0 pps
    Output packets:
                                  519498
                                                            0 pps
   IPV6 multicast statistics:
    Input bytes :
                                                            0 bps
                                  156500
    Output bytes :
                                16681118
                                                            0 bps
    Input packets:
                                    1818
                                                            0 pps
    Output packets:
                                  362633
                                                            0 pps
 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
 Oueue counters:
                       Queued packets Transmitted packets
                                                                Dropped packets
   0 best-effort
                               882558
                                                     882558
   1 expedited-fo
                                    0
                                                                               0
                                                          0
   2 assured-forw
                                    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 :
 Output bytes :
                            16681118
 Input packets:
 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 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__

SEE ALSO

show interfaces

show interfaces statistics

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

Initial Configuration of Ethernet Interfaces | 3

100GbE, 40GbE, and 10GbE | 426

Configuring Gigabit Ethernet Policers

IN THIS SECTION

- Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 256
- Configuring Gigabit Ethernet Policers | 258
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 266

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 on how to configure a two-color policer and tricolor policer.

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

For 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), 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 30 on page 256 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 30: 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

Table 30: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs (Continued)

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
MAG III		
MAC accounting	Yes	Yes
MAC filtering	Yes	Yes
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

Table 30: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs (Continued)

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
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

Configuring MAC Address Accounting | 23

Configuring a Policer Overhead

Configuring Gigabit Ethernet Policers

IN THIS SECTION

- Overview | 258
- Configuring a Policer | 259
- Specifying an Input Priority Map | 260
- Specifying an Output Priority Map | 261
- Applying a Policer | 262
- Configuring MAC Address Filtering | 263
- Example: Configuring Gigabit Ethernet Policers | 264

Overview

On Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can define rate limits for premium and aggregate traffic received on the interface. These policies 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 (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), 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.

NOTE: On MX Series 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.

For information on understanding Ethernet frame statistics, see the *MX Series Layer 2 Configuration Guide*.

Configuring 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:

```
burst size = bandwidth x 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.

Specifying 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 ieee802.1 premium 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). For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 264.

NOTE: On IQ2 and IQ2-E interfaces and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

Specifying 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 31 on page 261 shows the predefined forwarding classes and their associated queue assignments.

Table 31: Default Forwarding Classes

Forwarding Class Name	Queue
best-effort	Queue 0
expedited-forwarding	Queue 1
assured-forwarding	Queue 2
network-control	Queue 3

For more information about CoS forwarding classes, see the Junos OS Class of Service User Guide for Routing Devices. For a configuration example, see "Example: Configuring Gigabit Ethernet Policers" on page 264.

Applying a Policer

On all MX Series Router interfaces, Gigabit Ethernet IQ, IQ2, and IQ2-E 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), 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 (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

These policiers 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 or nnnn.nnn.nnn, 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.

NOTE: On untagged 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 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.

NOTE: 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.

NOTE: 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.

Configuring MAC Address Filtering

You cannot explicitly define traffic with specific source MAC addresses to be rejected; however, for Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet DPCs on MX Series routers, you can block all incoming packets that do not have a source address specified in the accept-source-mac statement. For more information about the accept-source-mac statement, see "Applying a Policer" on page 262.

To enable this blocking, include the source-filtering statement at the [edit interfaces *interface-name* gigether-options] hierarchy level:

```
[edit interfaces interface-name gigether-options]
source-filtering;
```

For more information about the source-filtering statement, see "Configuring MAC Address Filtering for Ethernet Interfaces" on page 20.

To accept traffic even though it does not have a source address specified in the accept-source-mac statement, include the no-source-filtering statement at the [edit interfaces *interface-name* gigether-options] hierarchy level:

```
[edit interfaces interface-name gigether-options]
no-source-filtering;
```

Example: Configuring Gigabit Ethernet Policers

IN THIS SECTION

- Example | 264
- Example Configuration | 265

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;
            }
        }
    }
}
```

RELATED DOCUMENTATION

Configuring MAC Address Accounting | 23

Configuring a Policer Overhead

Configuring Gigabit Ethernet Two-Color and Tricolor Policers

IN THIS SECTION

- Overview | 266
- Configuring a Policer | 268
- Applying a Policer | 269
- Example: Configuring and Applying a Policer | 269

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 258), configuring two-color policers and tricolor marking policers requires that you configure a firewall filter.

Configuring 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.

Applying 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 <code>[edit firewall]</code> 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.

Example: Configuring and Applying 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

Configuring MAC Address Accounting | 23

Configuring a Policer Overhead

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.

RELATED DOCUMENTATION

Configuring 10-Gigabit Ethernet PICs | 158

Configuring 40-Gigabit Ethernet PICs | 200

Gigabit Ethernet Autonegotiation

IN THIS SECTION

- Gigabit Ethernet Autonegotiation Overview | 271
- Configuring Gigabit Ethernet Autonegotiation | 272

Use this topic for information about how to configure Autonegotiation on Gigabit Ethernet Interfaces.

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.

NOTE:

- When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.
- On ACX Series Universal Metro Routers, when the autonegotiation is disabled, the speed has to be explicitly configured to 10–100 Mbps.

• On T4000 routers, the auto-negotiation command is ignored for interfaces other than Gigabit Ethernet.

SEE ALSO

Configuring 100-Gigabit Ethernet MICs/PICs | 203

Configuring Gigabit Ethernet Autonegotiation

IN THIS SECTION

- Configuring Gigabit Ethernet Autonegotiation with Remote Fault | 272
- Configuring Flow Control | 272
- Configuring Autonegotiation Speed on MX Series Routers | 272
- Displaying Autonegotiation Status | 273

Configuring 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>
```

Configuring Flow Control

To enable flow control, include the flow-control statement at the [edit interfaces ge-fpc/pic>/port gigether-options] hierarchy level. For more information, see "Configuring Flow Control" on page 18.

Configuring Autonegotiation Speed on MX Series Routers

MX Series 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.

MX Series 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, theauto-10m-100m option allows the fixed tri-speed port to auto negotiate with ports limited by 100m or 10mmaximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC on MX platform. This option does not support other MICs on MX platform.

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.

NOTE: Starting in Junos OS Release 14.2, on MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the auto-10m-100m option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link. However, when you view the interface speed of the DPC, using the show interfaces command, the value of the speed is not accurately displayed. For instance, if you configure the speed of the Tri-rate enhanced DPC, as 100Mbps on both sides of the link, the interface speed of the DPC is negotiated to 100 Mbps. However, the interface speed of the DPC displays 1 bps. This is an issue with the show interfaces command only. The actual interface speed is 100 Mbps.

Displaying Autonegotiation Status

To display Gigabit Ethernet interface details, including the autonegotiation status, use the operational mode command show interfaces ge- fpcl picl port extensive.

Table 32 on page 274 and Table 33 on page 277 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 32: 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	
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	

Table 32: Mode and Autonegotiation Status (Local) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
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	
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

Table 32: Mode and Autonegotiation Status (Local) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
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	
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	

Table 32: Mode and Autonegotiation Status (Local) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
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 33: 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	

Table 33: Mode and Autonegotiation Status (Remote) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
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	
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	

Table 33: Mode and Autonegotiation Status (Remote) (Continued)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
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	

RELATED DOCUMENTATION

Configuring 10-Gigabit Ethernet PICs | 158

Configuring Layer 2 Overhead Attribute in Interface Statistics | 247

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
14.2	Starting from Junos OS Release 14.2, theauto-10m-100m option allows the fixed tri-speed port to auto negotiate with ports limited by 100m or 10mmaximum speed.
14.2	Starting in Junos OS Release 14.2, on MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the auto-10m-100m option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link.

RELATED DOCUMENTATION

Configuring 40-Gigabit Ethernet PICs | 200

Configuring 100-Gigabit Ethernet MICs/PICs | 203



Port Speed

Port Speed Overview | 281

Configure Port Speed | 288

Port Speed on ACX Routers | 294

Port Speed on MX Routers | 321

Port Speed on PTX Routers | 374

Port Speed Overview

IN THIS SECTION

Introduction to Port Speed | 281

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 | 282
- What is Oversubscription? | 283
- Supported Line Cards and Devices | 283

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 34 on page 282 describes the different types of port speed configuration.

Table 35 on page 283 describes the naming formats for the channelized and non-channelized interfaces.

Table 36 on page 284 lists the line cards and devices, the port profile configuration, and link to the same.

"Configure Port Speed - PIC Level" on page 288 describes the port speed configuration at PIC Level - Chassis Hierarchy.

"Configure Port Speed - Port Level" on page 289 describes the port speed configuration at Port Level - Chassis Hierarchy.

"Configure Port Speed for Non-Channelized Interfaces" on page 291 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 292 describes the steps to configure the port speed for channelized interfaces from the [edit interfaces] hierarchy.

Table 34: 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 pic-mode 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 speed 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 35: 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.

Supported Line Cards and Devices

Table 36 on page 284 lists the line cards and devices, the port profile configuration, and link to the same.

Table 36: Supported Line Cards and Routers

Product Family	Line Card or Device	Supported Port Profile Configuration
ACX Series Routers	ACX7509-FPC-20Y	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces
	ACX7509-FPC-16C	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces
	ACX7509-FPC-4CD	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces
	ACX7100-48L	Interfaces Hierarchy Channelized Interfaces Non-channelized Interfaces
	ACX7100-32C	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces
	ACX710	Chassis Hierarchy Channelized Interfaces Non-channelized Interfaces

Table 36: Supported Line Cards and Routers (Continued)

Product Family	Line Card or Device	Supported Port Profile Configuration
	ACX7024	Chassis HierarchyPIC levelPort level
	ACX5448, ACX5448-D, and ACX5448-M	Chassis Hierarchy PIC level Port level
MX Series Routers	MPC7E-MRATE	Chassis Hierarchy PIC level Port level
	MPC7E-10G	Chassis Hierarchy PIC level Port level
	MIC-MRATE	Chassis Hierarchy PIC level Port level
	MX10003 MPC	Chassis Hierarchy PIC level Port level

Table 36: Supported Line Cards and Routers (Continued)

Product Family	Line Card or Device	Supported Port Profile Configuration
	MX204	Chassis Hierarchy PIC level Port level
	MX304	Chassis Hierarchy PIC level Port level
	MIC-MACSEC-20GE	Chassis Hierarchy PIC level Port level
	MPC10E-10C-MRATE	Chassis Hierarchy PIC level Port level
	MPC10E-15C-MRATE	Chassis Hierarchy PIC level Port level
	MX2K-MPC11E	Chassis Hierarchy PIC level Port level

Table 36: Supported Line Cards and Routers (Continued)

Product Family	Line Card or Device	Supported Port Profile Configuration
	MX10K-LC2101	Chassis Hierarchy PIC level Port level
	MX10K-LC9600	Chassis Hierarchy PIC level Port level
PTX Series Routers	PTX10K-LC1201	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces
	PTX10001-36MR	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces
	PTX10K-LC1202-36MR	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces
	PTX10003	Interfaces HierarchyChannelized InterfacesNon-channelized Interfaces

SEE ALSO

number-of	ports
number-of	sub-ports
pic-mode	
speed	
Unused	

Configure Port Speed

IN THIS SECTION

- Configure Port Speed PIC Level | 288
- Configure Port Speed Port Level | 289
- Configure Port Speed for Non-Channelized Interfaces | 291
- Configure Port Speed for Channelized Interfaces | 292

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 heirarchy, 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.

```
}
...
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 | 294
- Port Speed on ACX5448 Router | 296
- Port Speed on ACX7024 Router | 299
- Port Speed on ACX7100 Router | 301
- Port Speed on ACX7332 and ACX7348 Routers | 307
- Port Speed on ACX7509 Router | 312

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 37: 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 38: 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/ <i>0</i> - xe-0/0/ <i>23</i>
	1-Gigabit Ethernet interface (Tri-rate SFP-T optics)	xe-0/0/ <i>0</i> - xe-0/0/ <i>15</i>
	100-Mbps and 1-Gigabit Ethernet interface (Tri-rate SFP-T optics)	xe-0/0/ <i>16</i> - xe-0/0/ <i>23</i>

Table 38: Interface Naming Convention for ACX710 (Continued)

PIC	Interface type	Interfaces
PIC 1	100-Gigabit Ethernet interface (4 QSFP28 ports)	et-0/1/ <i>0</i> - et-0/1/ <i>3</i>

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 otinterface, 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 39: 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 39: 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 40: 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 41: 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	O 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 288 . For information on how to configure the speed at the port level, see "Configure Port Speed - Port Level" on page 289 .

Table 42: 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 42: 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 Etherner interface (CFP2-DCO ports)	Not Supported	ot-0/2/38 and ot-0/2/39	Not Supported

For information on multipelxing on ACX5448-D routers, see Multiplexing on ACX5448-D Routers.

Port Speed on ACX7024 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 routers, see ACX7024 Universal Metro Router Hardware Guide.

For information about platform support, see Hardware Compatibility Tool (HCT)

Table 43: Port speed for ACX7024 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 27 If you configure Precision Time Protocol (PTP), port 27 is disabled.	1-Gigabit Ethernet10-Gigabit Ethernet25-Gigabit EthernetDefault: 10-Gigabit Ethernet

Table 44: Interface Naming Convention for ACX7024 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 44: Interface Naming Convention for ACX7024 routers (Continued)

Optics Type	Interface type	Interfaces
		et-0/0/0:3 - et-0/0/3:3
	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 291. 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 292.

Port Speed on ACX7100 Router

IN THIS SECTION

- Port Speed on ACX7100-48L Router | 302
- Port Speed on ACX7100-32C Router | 305

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.

Table 45: 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 If you configure Precision Time Protocol (PTP), port 47 is disabled.	10-Gigabit Ethernet 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 46: 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 9216 bytes for transit traffic.
Forward Error Correction (FEC) support	By default, FEC119 (KP4) is enabled on 400-Gigabit Ethernet and 200-Gigabit Ethernet interfaces. 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. For information about FEC support, see fec (gigether).

Table 47: Interface Naming Convention for ACX7100-48L Router

PIC	Interface Type	Interfaces
PIC 0	10-Gigabit Ethernet interface25-Gigabit Ethernet interface50-Gigabit Ethernet interface	et-0/0/ <i>0</i> -et-0/0/ <i>47</i>

Table 47: Interface Naming Convention for ACX7100-48L Router (Continued)

400-Gigabit Ethernet 1x40-Gigabit Ethernet 1x100-Gigabit Ethernet 4x100-Gigabit Ethernet 2x100-Gigabit Ethernet 1x100-Gigabit Ethernet 8x50-Gigabit Ethernet	PIC	Interface Type	Interfaces
4x10-Gigabit Ethernet 4x25-Gigabit Ethernet interface (6 QSFP56-DD ports)		1x40-Gigabit Ethernet 1x100-Gigabit Ethernet 4x100-Gigabit Ethernet 2x100-Gigabit Ethernet 1x100-Gigabit Ethernet 8x50-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet	et-0/0/ <i>48</i> -et-0/0/ <i>53</i>

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 48: ACX7100-48L Port Macro Port Speed Configuration Router

Per Port Macro	Port Speed Configuration			Support ed
Port	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 Harware 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.

Table 49: 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 10-Gigabit Ethernet

Table 49: Port Speed for ACX7100-32C (Continued)

PIC	Port Number	Port Speed Supported
	Port 31	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 50: 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 51: 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/ <i>0</i> -et-0/0/ <i>31</i>
	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/ <i>32</i> -et-0/0/ <i>35</i>

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 291. 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 292.

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. QSPF 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 52: 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 53: 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

Table 53: Port Speed for ACX7348 Fixed FPC Interface (Continued)

Slot	Port Number	Supported Port Speed
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 54: 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	
1,2	1,3,4,5	4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 1x40-Gigabit Ethernet 1x100-Gigabit Ethernet	

Table 54: Port Speed for 6xQSFPDD FPC Line Card (ACX7332 and ACX7348) (Continued)

Slot	Port Number	Supported Port Speed
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 55: 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
1,2	(0,1,3) (2,4,5)	2x100G 4x25G 4x10G 1x100G 1x40G	4x25G 4x10G	Unused

Table 55: 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 4x10G 1x100G 1x40G	Unused	4x25G 4x10G 4x50G
3	(0,1,2)	1x100G 4x25G 1x40G 4x10G	1x100G 4x25G 1x40G 4x10G	Unused

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 56: 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.

Table 56: Port Speed for 16xSFP56 FPC Line Card (ACX7332 and ACX7348) (Continued)

Slot	Port Number	Supported Port Speed
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 291. 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 292.

Port Speed on ACX7509 Router

IN THIS SECTION

- Port Speed for ACX7509-FPC-20Y Line Card | 313
 - Port Speed on ACX7509-FPC-16C Line Card | 315
- Port Speed on ACX7509-FPC-4CD Line Card | 319

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 57: 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 58: 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

Table 58: Port Group for ACX7509-FPC-20Y Line Card (Continued)

Slot	Port Number	Supported Port Speed
1 and 5	(0-7), (8-16), and (16-19)	1x25-Gigabit Ethernet 1x10-Gigabit Ethernet

Table 59: ACX7509-FPC-20Y Line Card Details and Description

Details	Description
Autonegotation	Autonegotiation is supported only on the following channelization: • 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 9216 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 fec (gigether).

Table 60: Interface Naming Convention for ACX7509-FPC-20Y Line Card

PIC	Interface Type	Interfaces
PIC 0	1-Gigabit Ethernet 10-Gigabit Ethernet	et-0/0/0-et-0/0/19
	25-Gigabit Ethernet	et-0/0/0-et-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 61: 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
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

Table 61: Port Speed for ACX7509-FPC-16C Line Card (Continued)

Slot	Port Number	Supported Port Speed
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
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

Table 61: Port Speed for ACX7509-FPC-16C Line Card (Continued)

Slot	Port Number	Supported Port Speed
slots (0,2,3,4,6, and 7)	Ports: 2,3,6,7,10,11,14, and 15	Not supported

Table 62: ACX7509-FPC-16C Line Card Details and Description

Details	Description
Channelization	You can channelize the ports operating at: • 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 9216 bytes for transit traffic.
Forward Error Correction (FEC) support	FEC91 is enabled on 100-Gigabit and 25-Gigabit Ethernet interfaces. FEC91 is enabled by default on 25-Gigabit Ethernet interfaces. On the Direct Attach Copper (DAC) connection of the 25-Gigabit Ethernet interface, FEC74 is enabled. 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. For information about FEC support, see fec (gigether).

Table 63: 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
	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]

Table 63: 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-0/0/14	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 64: 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 65: ACX7509-FPC-4CD Line Card Details and Description

Details	Description
Channelization	You can channelize the ports operating at: • 400-Gbps speed to 4x100-Gbps speed.
Port speed configuration	Supports interfaces [edit interfaces interface-name] hierarchy.

Table 65: ACX7509-FPC-4CD Line Card Details and Description (Continued)

Details	Description
MTU size	ACX7509-FPC-4CD WAN interfaces support maximum MTU size of 9216 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 66: 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 291. 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 292.

RELATED DOCUMENTATION

request security system-keystore unlink

speed (Ethernet)

Port Speed on MX Routers

IN THIS SECTION

- Port Speed on MX204 | 321
- Port speed on MX304 Router | 326
- Port Speed on MPC7E-MRATE | 329
- Port Speed on MPC7E-10G | 332
- Port Speed on MIC-MRATE | 334
- Port Speed on MX10003 MPC | 338
- Port Speed on MIC-MACSEC-20GE | 344
- Port Speed on MPC10E-10C-MRATE | 345
- Port Speed on MPC10E-15C-MRATE | 350
- Port Speed on MX2K-MPC11E | 356
- Port Speed on MX10K-LC480 | 366
- Port Speed on MX10K-LC2101 | 368
- Port Speed on MX10K-LC9600 | 372

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 67: 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)
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 68: Maximum number of ports Configurable at PIC or Port level (MX204)

Type of Ports	PIC Mode (maximum configurable ports on PICO and PIC1)	Port Mode (maximum configurable ports on PICO and PIC1)
10/1 Gigabit Ethernet Ports	24 16 ports from PICO and 8 ports from PIC1	20 12 ports from PICO and 8 ports from PIC1

Table 68: Maximum number of ports Configurable at PIC or Port level (MX204) (Continued)

Type of Ports	PIC Mode (maximum configurable ports on PIC0 and PIC1)	Port Mode (maximum configurable ports on PICO and PIC1)
40 Gigabit Ethernet Ports	4 4 ports from PICO	4
100 Gigabit Ethernet Ports	4 4 ports from PICO	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 69: Valid Port Speed Combinations at Port Level (MX204)

Port Mode											
PICO			PIC1								
100	100	100	100	O Configure the number of active ports as O							
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
100	10/40	10/40	10/40	10 10 10 10 10 10 10							
10/40	10/40	10/40	10/40	O Configure the number of active ports as O							

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 70: Valid Port Speed Combinations at PIC level (MX204)

PIC Mode						
PICO (with four multir	PICO (with four multirate ports)					
100	100 100 100					
40	40	40	40	O Configure the number of active ports as 0		
10	10	10	10	O Configure the number of active		

Table 71: Interface Naming Conventions (MX204)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe- <i>x</i> /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]

Table 71: Interface Naming Conventions (MX204) (Continued)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
1	xe-x/1/0:[0-3]	et-x/1/0:[0-3]	et-x/1/0:[0-3]
	xe- <i>x</i> /1/1:[0-3]	xe- <i>x</i> /1/1:[0-3]	xe- <i>x</i> /1/1:[0-3]
	xe- <i>x</i> /1/2:[0-3]	xe- <i>x</i> /1/2:[0-3]	xe- <i>x</i> /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- <i>x</i> /1/4:[0-3]	xe-x/1/4:[0-3]
	xe-x/1/5:[0-3]	xe- <i>x</i> /1/5:[0-3]	xe- <i>x</i> /1/5:[0-3]
	xe-x/1/6:[0-3]	xe-x/1/6:[0-3]	xe-x/1/6:[0-3]
	xe- <i>x</i> /1/7:[0-3]	xe- <i>x</i> /1/7:[0-3]	xe-x/1/7:[0-3]

Table 72: Active Physical Ports (MX204) at PIC Level

PIC	Number of Ports Configured	Active Ports		
PIC0		10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
	0	-	-	-
	1	0	0	0
	2	0, 1	0, 1	0, 1

Table 72: Active Physical Ports (MX204) at PIC Level (Continued)

PIC	Number of Ports Configured	Active Ports		
	3	0, 1, 2	0, 1, 2	0, 1, 2
PIC1	0	-	-	-
	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	-	-

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.

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 73: 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 73: 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 1x40-Gigabit Ethernet Supports 1-Gbps speed (Junos OS Release 23.1R1 onwards)

Table 74: Interface Naming Convention for MX304

PIC	Interface Type	Interfaces
PIC 0	1x1-Gigabit Ethernet Interface	ge-0/0/0

Table 74: 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 1x100-Gigabit Ethernet Interface	et-0/0/0
	1x400-Gigabit Ethernet Interface	
	4x25-Gigabit Ethernet Interface 4x100-Gigabit Ethernet Interface	et-0/0/0:0 - et-0/0/3:0
	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

Port Speed on MPC7E-MRATE

All the ports in the MPC7E-MRATE MPC operate as four 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 75: Port Speed for MPC7E-MRATE

PIC	Port	Port Speed Supported
PIC 0 (Fixed-port PIC)	0, 1, 3, 4 2 and 5	40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC 1 (Fixed-port PIC)	0, 1, 3,4 2 and 5	40-Gigabit Ethernet 100-Gigabit Ethernet Default: 10-Gigabit Ethernet

Table 76 on page 330 lists the naming conventions used for interfaces on MPC7E-MRATE for MX240, MX480, MX960, MX2010, and MX2020 routers.

Table 76: Interface Naming Convention for MPC7E-MRATE

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0	-
	xe- <i>x</i> /0/1: <i>[0-3]</i>	et- <i>x</i> /0/1	-
	xe- <i>x</i> /0/2: <i>[0-3]</i>	et- <i>x</i> /0/2	et- <i>x</i> /0/2
	xe-x/0/3:[0-3]	et-x/0/3	-
1	xe- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4	-
	xe- <i>x</i> /0/5: <i>[0-3]</i>	et- <i>x</i> /0/5	et- <i>x</i> /0/5
	xe- <i>x</i> /1/0: <i>[0-3]</i>	et- <i>x</i> /1/0	-

Table 76: Interface Naming Convention for MPC7E-MRATE (Continued)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
	xe- <i>x</i> /1/1: <i>[0-3]</i>	et- <i>x</i> /1/1	-
	xe- <i>x</i> /1/2: <i>[0-3]</i>	et- <i>x</i> /1/2	et- <i>x</i> /1/2
	xe- <i>x</i> /1/3: <i>[0-3]</i>	et- <i>x</i> /1/3	-
	xe- <i>x</i> /1/4: <i>[0-3]</i>	et- <i>x</i> /1/4	-
	xe- <i>x</i> /1/5: <i>[0-3]</i>	et- <i>x</i> /1/5	et- <i>x</i> /1/5

Table 77: Active Physical Ports on MPC7E-MRATE

Number of Active ports	Active Physical Ports for Different Configured Speeds		
Configured	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

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.

All the ports in the MPC7E-10G MPC operate as four 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 78: Port Speed for MPC7E-10G

PIC	Port	Port Speed Supported
PIC 0	20 (0-19)	10-Gigabit Ethernet Default: 10-Gigabit Ethernet
PIC 1	20 (0-19) 2 and 5	10-Gigabit Ethernet Default: 10-Gigabit Ethernet

Table 79 on page 332 lists the naming conventions used for interfaces on MPC7E-10G for MX240, MX480, MX960, MX2010, and MX2020 routers.

Table 79: Interface Naming Convention for MPC7E-10G

Packet Forwarding Engine	10-Gigabit Ethernet Interface
0	xe-x/0/0:[0-3]to xe-x/0/19:[0-3]
1	xe-x/1/0:[0-3]to xe-x/1/19: [0-3]

Table 80: 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 80: 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

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 81: 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 82 on page 335 lists the naming conventions used for interfaces on MIC-MRATE installed on Slot 0 of MPC8E and MPC9E.

Table 82: 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-x0/5:[0-3]	et- <i>x</i> /0/0: <i>[0-3]</i> -xe- <i>x</i> 0/5: <i>[0-3]</i>	et-x/0/0:[0-3]-xe-x0/3:[0-3]
1	xe- <i>x</i> /0/6: <i>[0-3]</i> -xe- <i>x</i> /0/11: <i>[0-3]</i>	et- <i>x</i> /0/6: <i>[0-3]</i> -xe- <i>x</i> 0/11: <i>[0-3]</i>	et- <i>x</i> /0/6: <i>[0-3]</i> -xe- <i>x</i> 0/11: <i>[0-3]</i>

Table 83 on page 335 lists the naming conventions used for interfaces on MIC-MRATE installed on Slot 1 of MX10003 MPC.

Table 83: 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:	et-x/0/0:[0-3]to xe-x/0/3:	et- <i>x</i> /0/0: <i>[0-3]</i> to xe- <i>x</i> /0/3:
	[0-3]	[0-3]	<i>[0-3]</i>

Table 83: 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:	et-x/0/4:[0-3]to xe-x/0/7:	et- <i>x</i> /0/6: <i>[0-3]</i> to xe- <i>x</i> /0/7:
	[0-3]	[0-3]	<i>[0-3]</i>
2	xe-x/0/8:[0-3]to xe-x/0/11:	xe-x/0/8:[0-3]toxe-x/0/11:	xe- <i>x</i> /0/8: <i>[0-3]</i> to xe- <i>x</i> /0/9:
	[0-3]	[0-3]	<i>[0-3]</i>

Table 84: Active Physical Ports on MIC-MRATE (MPC8E)

Number of Active Ports	Active Physical Ports for Different Configured Speeds			
Configured	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 84: Active Physical Ports on MIC-MRATE (MPC8E) (Continued)

Number of Active Ports	Active Physical Ports for Different Configured Speeds		
Configured	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 85 on page 337 lists the active physical ports on MIC-MRATE (MPC9E and MPC8E (1.6T))

Table 85: Active Physical Ports on MIC-MRATE (MPC9E and MPC8E (1.6T mode))

Number of Active Ports	Active Physical Ports for Different Configured Speeds		
Configured	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 85: Active Physical Ports on MIC-MRATE (MPC9E and MPC8E (1.6T mode)) (Continued)

Number of Active Ports	Active Physical Ports for Different	t Configured Speeds	
Configured	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

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 86: 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-MASCEC MIC, all the 4x10-Gigabit Ethernet ports can operate in 1 Gbps mode.

Table 87: PFE Based Port Mode Configuration

Port Speed configuration on PIC1(Gbps)				Port speed configuration on PICO(Gbps)	
100	100	100	100	0	0
100	100	100	10/40	0	0
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

Table 87: PFE Based Port Mode Configuration (Continued)

Port Speed configuration on PIC1(Gbps)				Port speed configuration or	PICO(Gbps)
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 88: PIC Mode Configuration

Port Speed configuration on PIC1(Gbps)	Port speed configuration on PICO(Gbps)
100	O Configure the number of ports to 0.
10	10
40	40

Table 89 on page 340 lists the naming conventions used for interfaces on the fixed-port PIC when installed in slot 0 of the MX10003 MPC. Table 90 on page 341 lists the naming conventions used for interfaces on the modular MIC when installed in slot 1 of the MPC.

Table 89: 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- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0	-
	xe- <i>x</i> /0/1: <i>[0-3]</i>	et- <i>x</i> /0/1	-

Table 89: Interface Naming Convention for the Fixed-Port PIC in Slot 0 of MX10003 MPC (Continued)

Packet Forwarding Engine	10-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
1	xe- <i>x</i> /0/2: <i>[0-3]</i>	et- <i>x</i> /0/2	-
	xe- <i>x</i> /0/3: <i>[0-3]</i>	et- <i>x</i> /0/3	-
2	xe- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4	-
	xe- <i>x</i> /0/5: <i>[0-3]</i>	et- <i>x</i> /0/5	-

Table 90: 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- <i>x</i> /1/0: <i>[0-3]</i>	et- <i>x</i> /1/0	et- <i>x</i> /1/0
	xe- <i>x</i> /1/1: <i>[0-3]</i>	et- <i>x</i> /1/1	et- <i>x</i> /1/1
	xe- <i>x</i> /1/2: <i>[0-3]</i>	et- <i>x</i> /1/2	et- <i>x</i> /1/2
	xe- <i>x</i> /1/3: <i>[0-3]</i>	et- <i>x</i> /1/3	et- <i>x</i> /1/3
1	xe- <i>x</i> /1/4: <i>[0-3]</i>	et- <i>x</i> /1/4	et- <i>x</i> /1/4
	xe- <i>x</i> /1/5 <i>[0-3]</i>	et- <i>x</i> /1/5	et- <i>x</i> /1/5
	xe- <i>x</i> /1/6 <i>[0-3]</i>	et- <i>x</i> /1/6	et- <i>x</i> /1/6
	xe- <i>x</i> /1/7: <i>[0-3]</i>	et- <i>x</i> /1/7	et- <i>x</i> /1/7
2	xe- <i>x</i> /1/8: <i>[0-3]</i>	et- <i>x</i> /1/8	et- <i>x</i> /1/8

Table 90: 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- <i>x</i> /1/9: <i>[0-3]</i>	et- <i>x</i> /1/9	et- <i>x</i> /1/9
	xe- <i>x</i> /1/10: <i>[0-3]</i>	et- <i>x</i> /1/10	et- <i>x</i> /1/10
	xe- <i>x</i> /1/11: <i>[0-3]</i>	et- <i>x</i> /1/11	et- <i>x</i> /1/11

Table 91: 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	
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	

Table 91: 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	
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 92: Active Physical Ports on MX10003 MPC at the PIC level

Ports Configured (number-of-ports	Active Physical Ports for Different Configured Speeds		
Statement)	10-Gigabit Ethernet	40-Gigabit Ethernet	
1	0	0	
2	0, 2	0, 2	
3	0, 2, 4	0, 2, 4	
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	

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 autonegotiation and set the speed to 1-Gbps on the peer side to bring the link up on the peer side.

Table 93: 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

Table 93: Port Speed for MIC-MACSEC-20GE (Continued)

PIC	Port Number	Port Speed Supported
PIC 1	8-9	10-Gigabit Ethernet Default: 10-Gigabit Ethernet

Table 94: Interface Naming Conventions (MIC-MACSEC-20GE)

PIC	1-Gigabit Ethernet Interface	10-Gigabit Ethernet Interface
PIC 0	ge-x/0/[0-9]	xe- <i>x</i> /0/0
PIC 1	ge-x/1/[0-9]	xe- <i>x</i> /1/0
PIC 2	ge-x/2/[0-9]	xe- <i>x</i> /2/0
PIC 3	ge-x/3/[0-3]	xe- <i>x</i> /3/0

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 95 on page 346 summarizes the Packet Forwarding Engine mapping and the supported port speeds.

Table 95: 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.

Table 96: 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 97: 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 97: 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 98: 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- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0	et- <i>x</i> /0/0	-
	xe- <i>x</i> /0/1: <i>[0-3]</i>	et- <i>x</i> /0/1: <i>[0-3]</i>	et- <i>x</i> /0/1	et- <i>x</i> /0/1	-
	xe- <i>x</i> /0/2: <i>[0-3]</i>	et- <i>x</i> /0/2: <i>[0-3]</i>	et- <i>x</i> /0/2	et- <i>x</i> /0/2	-
	xe- <i>x</i> /0/3: <i>[0-3]</i>	et- <i>x</i> /0/3: <i>[0-3]</i>	et- <i>x</i> /0/3	et- <i>x</i> /0/3	-
	xe- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4	et- <i>x</i> /0/4	et- <i>x</i> /0/4
1	xe- <i>x</i> /1/0: <i>[0-3]</i>	et- <i>x</i> /1/0: <i>[0-3]</i>	et- <i>x</i> /1/0	et- <i>x</i> /1/0	-
	xe- <i>x</i> /1/1: <i>[0-3]</i>	et- <i>x</i> /1/1: <i>[0-3]</i>	et- <i>x</i> /1/1	et- <i>x</i> /1/1	-
	xe- <i>x</i> /1/2: <i>[0-3]</i>	et- <i>x</i> /1/2: <i>[0-3]</i>	et- <i>x</i> /1/2	et- <i>x</i> /1/2	-

Table 98: 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- <i>x</i> /1/3: <i>[0-3]</i>	et- <i>x</i> /1/3: <i>[0-3]</i>	et- <i>x</i> /1/3	et- <i>x</i> /1/3	-
	xe- <i>x</i> /1/4: <i>[0-3]</i>	et- <i>x</i> /1/4: <i>[0-3]</i>	et- <i>x</i> /1/4	et- <i>x</i> /1/4	et- <i>x</i> /0/4

Table 99: Active Ports with port speed configured at PIC level

PIC Type	Num ber of Activ e Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSF P28	1	0	0	0	0	4
PIC (PIC 0)	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,	0,1,2,3,4	0, 1, 2, 3, 4	0, 1, 2, 3, 4	4
5xQSF P28	1	0	0	0	0	4
PIC (PIC 1)	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 99: Active Ports with port speed configured at PIC level (Continued)

PIC Type	Num ber of Activ e Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
	5	0, 1, 2, 3,	0,1,2,3,4	0, 1, 2, 3, 4	0, 1, 2, 3, 4	4

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

To view the supported transceivers, optical interfaces, and DAC cables on MPC10E-15C-MRATE, see Hardware Compatibility Tool (HCT).

Table 100: 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.

Table 101: Port speed support on MPC10E-15C-MRATE

PIC	PIC leve	l 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 102: 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
	400GE + 4x25GE	Yes

Table 102: Port Speed Combinations with 400GE on MPC10E-15C-MRATE (Continued)

PIC	Port Speed Combinations	Supported
	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 103: 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- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0	et- <i>x</i> /0/0	-
	xe- <i>x</i> /0/1: <i>[0-3]</i>	et- <i>x</i> /0/1: <i>[0-3]</i>	et- <i>x</i> /0/1	et- <i>x</i> /0/1	-
	xe- <i>x</i> /0/2: <i>[0-3]</i>	et- <i>x</i> /0/2: <i>[0-3]</i>	et-x/0/2	et-x/0/2	-
	xe- <i>x</i> /0/3: <i>[0-3]</i>	et- <i>x</i> /0/3: <i>[0-3]</i>	et-x/0/3	et-x/0/3	-
	xe- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4	et-x/0/4	et- <i>x</i> /0/4
1	xe- <i>x</i> /1/0: <i>[0-3]</i>	et- <i>x</i> /1/0: <i>[0-3]</i>	et- <i>x</i> /1/0	et- <i>x</i> /1/0	-
	xe- <i>x</i> /1/1: <i>[0-3]</i>	et- <i>x</i> /1/1: <i>[0-3]</i>	et- <i>x</i> /1/1	et- <i>x</i> /1/1	-
	xe- <i>x</i> /1/2: <i>[0-3]</i>	et- <i>x</i> /1/2: <i>[0-3]</i>	et- <i>x</i> /1/2	et- <i>x</i> /1/2	-
	xe- <i>x</i> /1/3: <i>[0-3]</i>	et- <i>x</i> /1/3: <i>[0-3]</i>	et- <i>x</i> /1/3	et- <i>x</i> /1/3	-
	xe- <i>x</i> /1/4: <i>[0-3]</i>	et- <i>x</i> /1/4: <i>[0-3]</i>	et- <i>x</i> /1/4	et- <i>x</i> /1/4	et- <i>x</i> /1/4
2	xe- <i>x</i> /2/0: <i>[0-3]</i>	et- <i>x</i> /2/0: <i>[0-3]</i>	et- <i>x</i> /2/0	et- <i>x</i> /2/0	-
	xe- <i>x</i> /2/1: <i>[0-3]</i>	et- <i>x</i> /2/1: <i>[0-3]</i>	et- <i>x</i> /2/1	et- <i>x</i> /2/1	-
	xe- <i>x</i> /2/2: <i>[0-3]</i>	et- <i>x</i> /2/2: <i>[0-3]</i>	et- <i>x</i> /2/2	et- <i>x</i> /2/2	-
	xe- <i>x</i> /2/3: <i>[0-3]</i>	et- <i>x</i> /2/3: <i>[0-3]</i>	et- <i>x</i> /2/3	et- <i>x</i> /2/3	-

Table 103: 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- <i>x</i> /2/4: <i>[0-3]</i>	et- <i>x</i> /2/4: <i>[0-3]</i>	et- <i>x</i> /2/4	et- <i>x</i> /2/4	et- <i>x</i> /2/4

Table 104: Active Ports with port speed configuration at PIC level

PIC Type	Numbe r of Active Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
5xQSFP 28 PIC	1	0	0	0	0	4
(PIC 0)	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	1	0	0	0	0	4
(PIC 1)	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	1	0	0	0	0	4

Table 104: Active Ports with port speed configuration at PIC level (Continued)

PIC Type	Numbe r of Active Ports	10-Gigabit Ethernet	25-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet	400-Gigabit Ethernet
(PIC 2)	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

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 105: Port Speed for the MX2K-MPC11E

PIC	Port Number	Port Speed Supported		
PIC 0 (or PFE 0)	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 1 (or PFE 1)	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 105: 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 105: 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 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 5 (or PFE 5)	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 105: 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 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 7 (or PFE 7)	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 106: 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- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0	et-x/0/0	et- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0
			et- <i>x</i> /0/1		-
			et-x/0/2		-
			et- <i>x</i> /0/3		-
			et- <i>x</i> /0/4		-
1	xe- <i>x</i> /1/0: <i>[0-3]</i>	et- <i>x</i> /1/0	et- <i>x</i> /1/0	et- <i>x</i> /1/0: <i>[0-3]</i>	et-x/1/0
			et- <i>x</i> /1/1		-
			et- <i>x</i> /1/2		-
			et- <i>x</i> /1/3		-
			et- <i>x</i> /1/4		-
2	xe- <i>x</i> /2/0: <i>[0-3]</i>	et- <i>x</i> /2/0	et- <i>x</i> /2/0	et- <i>x</i> /2/0: <i>[0-3]</i>	et-x/2/0
	[0 3]		et- <i>x</i> /2/1		-
			et- <i>x</i> /2/2		-
			et- <i>x</i> /2/3		-

Table 106: 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- <i>x</i> /2/4		-
3	xe- <i>x</i> /3/0: <i>[0-3]</i>	et- <i>x</i> /3/0	et-x/3/0	et- <i>x</i> /3/0: <i>[0-3]</i>	et-x/3/0
			et- <i>x</i> /3/1		-
			et-x/3/2		-
			et-x/3/3		-
			et- <i>x</i> /3/4		-
4	xe- <i>x</i> /4/0: <i>[0-3]</i>	et-x/4/0	et- <i>x</i> /4/0	et- <i>x</i> /4/0: <i>[0-3]</i>	et-x/4/0
			et- <i>x</i> /4/1		-
			et- <i>x</i> /4/2		-
			et-x/4/3		-
			et- <i>x</i> /4/4		-
5	xe- <i>x</i> /5/0: <i>[0-3]</i>	et- <i>x</i> /5/0	et- <i>x</i> /5/0	et- <i>x</i> /5/0: <i>[0-3]</i>	et-x/5/0
			et- <i>x</i> /5/1		-
			et- <i>x</i> /5/2		-

Table 106: 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- <i>x</i> /5/4		-
6	xe- <i>x</i> /6/0: <i>[0-3]</i>	et- <i>x</i> /6/0	et- <i>x</i> /6/0	et- <i>x</i> /6/0: <i>[0-3]</i>	et-x/6/0
			et- <i>x</i> /6/1		-
			et- <i>x</i> /6/2		-
			et- <i>x</i> /6/3		-
			et- <i>x</i> /6/4		-
7	xe- <i>x</i> /7/0: <i>[0-3]</i>	et- <i>x</i> /7/0	et- <i>x</i> /7/0	et- <i>x</i> /7/0: <i>[0-3]</i>	et- <i>x</i> /7/0
			et- <i>x</i> /7/1		-
			et- <i>x</i> /7/2		-
			et- <i>x</i> /7/3		-
			et- <i>x</i> /7/4		-

Table 107: 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	О	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 107: 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 107: 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

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 108: 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 109: 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 110: 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 110: 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

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.

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 (QSFPP-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 111: Port Speed for MX10K-LC2101

PIC	Ports	Port Speed Supported
PICO	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 111: 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 112: Interface Naming Conventions (MX10K-LC2101)

PFE	10-Gigabit Ethernet Interface	40-Gigabit Ethernet	100-Gigabit Ethernet
0	xe-x/0/0:[0-3]	et- <i>x</i> /0/0	et- <i>x</i> /0/0
	xe-x/0/1:[0-3]	et- <i>x</i> /0/1	et- <i>x</i> /0/1
	xe-x/0/2:[0-3]	et- <i>x</i> /0/2	et- <i>x</i> /0/2
	xe-x/0/3:[0-3]	et- <i>x</i> /0/3	et- <i>x</i> /0/3
1	xe- <i>x</i> /1/0:[0-3]	et- <i>x</i> /1/0	et- <i>x</i> /1/0
	xe-x/1/1:[0-3]	et- <i>x</i> /1/1	et- <i>x</i> /1/1
	xe- <i>x</i> /1/2:[0-3]	et- <i>x</i> /1/2	et- <i>x</i> /1/2
	xe- <i>x</i> /1/3:[0-3]	et- <i>x</i> /1/3	et- <i>x</i> /1/3

Table 112: 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- <i>x</i> /2/0
	xe-x/2/1:[0-3]	et- <i>x</i> /2/1	et- <i>x</i> /2/1
	xe-x/2/2:[0-3]	et- <i>x</i> /2/2	et- <i>x</i> /2/2
	xe-x/2/3:[0-3]	et-x/2/3	et- <i>x</i> /2/3
3	xe-x/3/0:[0-3]	et-x/3/0	et-x/3/0
	xe-x/3/1:[0-3]	et- <i>x</i> /3/1	et- <i>x</i> /3/1
	xe-x/3/2:[0-3]	et- <i>x</i> /3/2	et- <i>x</i> /3/2
	xe-x/3/3:[0-3]	et- <i>x</i> /3/3	et- <i>x</i> /3/3
4	xe-x/4/0:[0-3]	et-x/4/0	et- <i>x</i> /4/0
	xe-x/4/1:[0-3]	et- <i>x</i> /4/1	et- <i>x</i> /4/1
	xe-x/4/2:[0-3]	et-x/4/2	et- <i>x</i> /4/2
	xe-x/4/3:[0-3]	et-x/4/3	et- <i>x</i> /4/3
5	xe- <i>x</i> /5/0:[0-3]	et- <i>x</i> /5/0	et- <i>x</i> /5/0
	xe-x/5/1:[0-3]	et- <i>x</i> /5/1	et- <i>x</i> /5/1
	xe- <i>x</i> /5/2:[0-3]	et- <i>x</i> /5/2	et- <i>x</i> /5/2

Table 112: Interface Naming Conventions (MX10K-LC2101) (Continued)

PFE	10-Gigabit Ethernet Interface	40-Gigabit Ethernet	100-Gigabit Ethernet
	xe- <i>x</i> /5/3:[0-3]	et- <i>x</i> /5/3	et- <i>x</i> /5/3

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 113: 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 113: MX10K-LC9600 - Software Support and Description (Continued)

Software Feature Support	Description
Forward Error Correction (FEC) support	By default, 200-Gigabit Ethernet interfaces and 400-Gigabit Ethernet interfaces support FEC119 (KP4) mode. 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. For information about FEC support, see fec (gigether).

Table 114 on page 373 summarizes the supported port speeds on MX10K-LC9600.

Table 114: Port Speed for MX10K-LC9600

PIC	Port Number	Port Speed Supported
PIC 0 through PIC 5	0-3	4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 2x100-Gigabit Ethernet 4x100-Gigabit Ethernet Default: All active ports operate in 400-Gigabit Ethernet mode.

Table 115: 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 115: Interface Naming Convention for MX10K-LC9600 (Continued)

PIC	Interface Type	Interfaces
	4x25-Gigabit Ethernet Interface 4x100-Gigabit Ethernet Interface	et-0/0/0:0 - et-0/0/0:3
	40-Gigabit Ethernet Interface 100-Gigabit Ethernet Interface 400-Gigabit Ethernet Interface	et-0/0/0
	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

Port Speed on PTX Routers

IN THIS SECTION

- Port Speed on PTX10K-LC1201 | **375**
- Port Speed on PTX10001-36MR Router | 383
- Port Speed on PTX10K-LC1202-36MR | 394
- Port Speed on PTX10003 | 402

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 <math>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 <math>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 116 on page 375 summarizes the line card details and their description.

Table 116: PTX10K-LC1201 Software Support and Description

Software Feature Support	Description				
Forward Error Correction (FEC)	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 116: PTX10K-LC1201 Software Support and Description (Continued)

Software Feature Support	Description
Port Speed Configuration	 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 117: 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 118: 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- <i>x</i> /0/0	et- <i>x</i> /0/0: <i>[0-7]</i>	et-x/0/0 et-x/0/0:[0-1] et-x/0/0:[0-3]	et-x/0/0: [0-2]	et- <i>x</i> /0/0
	et- <i>x</i> /0/1: [<i>0-7</i>]	et- <i>x</i> /0/1: <i>[0-7]</i>	et- <i>x</i> /0/1	et- <i>x</i> /0/1: <i>[0-7]</i>	et-x/0/1 et-x/0/1:[0-1] et-x/0/1:[0-3]	et- <i>x</i> /0/1: <i>[0-2]</i>	et- <i>x</i> /0/1
	et-x/0/2: [0-7]	et-x/0/2: [0-7]	et- <i>x</i> /0/2	et- <i>x</i> /0/2: [0-7]	et-x/0/2 et-x/0/2:[0-1] et-x/0/2:[0-3]	et- <i>x</i> /0/2: [0-2]	et- <i>x</i> /0/2
	et-x/0/3: [0-7]	et-x/0/3: [0-7]	et- <i>x</i> /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- <i>x</i> /0/3
	et- <i>x</i> /0/4: <i>[0-7]</i>	et- <i>x</i> /0/4: <i>[0-7]</i>	et- <i>x</i> /0/4	et- <i>x</i> /0/4: <i>[0-7]</i>	et-x/0/4 et-x/0/4:[0-1] et-x/0/4:[0-3]	et- <i>x</i> /0/4: <i>[0-2]</i>	et- <i>x</i> /0/4

Table 118: 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- <i>x</i> /0/5: [0-7]	et- <i>x</i> /0/5: <i>[0-7]</i>	et- <i>x</i> /0/5	et- <i>x</i> /0/5: <i>[0-7]</i>	et-x/0/5: [0-1] et-x/0/5: [0-3]	et-x/0/5: [0-2]	et- <i>x</i> /0/5
	et-x/0/6: [0-7]	et-x/0/6: [0-7]	et- <i>x</i> /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- <i>x</i> /0/6
	et-x/0/7: [0-7]	et- <i>x</i> /0/7: [<i>0-7</i>]	et- <i>x</i> /0/7	et- <i>x</i> /0/7: [<i>0-7</i>]	et-x/0/7 et-x/0/7:[0-1] et-x/0/7:[0-3]	et-x/0/7: [0-2]	et- <i>x</i> /0/7
	et-x/0/8: [0-7]	et- <i>x</i> /0/8: [<i>0-7</i>]	et- <i>x</i> /0/8	et- <i>x</i> /0/8: <i>[0-7]</i>	et-x/0/8 et-x/0/8:[0-1] et-x/0/8:[0-3]	et-x/0/8: [0-2]	et- <i>x</i> /0/8
	et- <i>x</i> /0/9: [0-7]	et-x/0/9: [0-7]	et- <i>x</i> /0/9	et- <i>x</i> /0/9: <i>[0-7]</i>	et-x/0/9 et-x/0/9:[0-1] et-x/0/9:[0-3]	et- <i>x</i> /0/9: [<i>0-2</i>]	et- <i>x</i> /0/9
	et-x/0/10: [0-7]	et- <i>x</i> /0/10: <i>[0-7]</i>	et- <i>x</i> /0/10	et- <i>x</i> /0/10: <i>[0-7]</i>	et-x/0/10 et-x/0/10:[0-1] et-x/0/10:[0-3]	et- <i>x</i> /0/10: <i>[0-2]</i>	et- <i>x</i> /0/10

Table 118: 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- <i>x</i> /0/11: <i>[0-7]</i>	et- <i>x</i> /0/11: <i>[0-7]</i>	et- <i>x</i> /0/11	et- <i>x</i> /0/11: <i>[0-7]</i>	et-x/0/11 et-x/0/11:[0-1] et-x/0/11:[0-3]	et-x/0/11: [0-2]	et- <i>x</i> /0/11
	et-x/0/12: [0-7]	et- <i>x</i> /0/12: <i>[0-7]</i>	et- <i>x</i> /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- <i>x</i> /0/12
	et-x/0/13: [0-7]	et- <i>x</i> /0/13: <i>[0-7]</i>	et- <i>x</i> /0/13	et- <i>x</i> /0/13: <i>[0-7]</i>	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- <i>x</i> /0/14: [0-7]	et- <i>x</i> /0/14: <i>[0-7]</i>	et- <i>x</i> /0/14	et- <i>x</i> /0/14: <i>[0-7]</i>	et-x/0/14 et-x/0/14:[0-1] et-x/0/14:[0-3]	et-x/0/14: [0-2]	et- <i>x</i> /0/14
	et-x/0/15:	et- <i>x</i> /0/15: <i>[0-7]</i>	et- <i>x</i> /0/15	et- <i>x</i> /0/15: <i>[0-7]</i>	et-x/0/15 et-x/0/15:[0-1] et-x/0/15:[0-3]	et- <i>x</i> /0/15: <i>[0-2]</i>	et- <i>x</i> /0/15
	et- <i>x</i> /0/16: <i>[0-7]</i>	et-x/0/16: [0-7]	et- <i>x</i> /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- <i>x</i> /0/16

Table 118: 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- <i>x</i> /0/17:	et- <i>x</i> /0/17: <i>[0-7]</i>	et- <i>x</i> /0/17	et- <i>x</i> /0/17: <i>[0-7]</i>	et-x/0/17 et-x/0/17:[0-1] et-x/0/17:[0-3]	et-x/0/17: [0-2]	et- <i>x</i> /0/17
	et- <i>x</i> /0/18:	et-x/0/18: [0-7]	et- <i>x</i> /0/18	et- <i>x</i> /0/18: <i>[0-7]</i>	et-x/0/18 et-x/0/18:[0-1] et-x/0/18:[0-3]	et- <i>x</i> /0/18: <i>[0-2]</i>	et- <i>x</i> /0/18
	et-x/0/19: [0-7]	et- <i>x</i> /0/19: <i>[0-7]</i>	et- <i>x</i> /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- <i>x</i> /0/19
	et-x/0/20: [0-7]	et-x/0/20: [0-7]	et- <i>x</i> /0/20	et- <i>x</i> /0/20: <i>[0-7]</i>	et-x/0/20 et-x/0/20:[0-1] et-x/0/20:[0-3]	et-x/0/20: [0-2]	et- <i>x</i> /0/20
	et- <i>x</i> /0/21: [0-7]	et-x/0/21: [0-7]	et- <i>x</i> /0/21	et- <i>x</i> /0/21: <i>[0-7]</i>	et-x/0/21 et-x/0/21:[0-1] et-x/0/21:[0-3]	et- <i>x</i> /0/21: <i>[0-2]</i>	et- <i>x</i> /0/21
	et- <i>x</i> /0/22: [<i>0-7</i>]	et- <i>x</i> /0/22: [<i>0-7</i>]	et- <i>x</i> /0/22	et- <i>x</i> /0/22: <i>[0-7]</i>	et-x/0/22 et-x/0/22:[0-1] et-x/0/22:[0-3]	et- <i>x</i> /0/22: <i>[0-2]</i>	et- <i>x</i> /0/22

Table 118: 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- <i>x</i> /0/23: <i>[0-7]</i>	et- <i>x</i> /0/23	et- <i>x</i> /0/23: <i>[0-7]</i>	et-x/0/23 et-x/0/23:[0-1] et-x/0/23:[0-3]	et-x/0/23: [0-2]	et- <i>x</i> /0/23
	et-x/0/24: [0-7]	et- <i>x</i> /0/24: <i>[0-7]</i>	et- <i>x</i> /0/24	et- <i>x</i> /0/24: <i>[0-7]</i>	et-x/0/24 et-x/0/24:[0-1] et-x/0/24:[0-3]	et- <i>x</i> /0/24: <i>[0-2]</i>	et- <i>x</i> /0/24
	et-x/0/25: [0-7]	et-x/0/25: [0-7]	et- <i>x</i> /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- <i>x</i> /0/26: <i>[0-7]</i>	et- <i>x</i> /0/26	et- <i>x</i> /0/26: <i>[0-7]</i>	et-x/0/26 et-x/0/26:[0-1] et-x/0/26:[0-3]	et-x/0/26: [0-2]	et- <i>x</i> /0/26
	et-x/0/27: [0-7]	et- <i>x</i> /0/27: <i>[0-7]</i>	et- <i>x</i> /0/27	et- <i>x</i> /0/27: <i>[0-7]</i>	et-x/0/27 et-x/0/27:[0-1] et-x/0/27:[0-3]	et- <i>x</i> /0/27: [<i>0-2</i>]	et- <i>x</i> /0/27
	et-x/0/28: [0-7]	et-x/0/28: [0-7]	et- <i>x</i> /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 118: 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- <i>x</i> /0/29: <i>[0-7]</i>	et- <i>x</i> /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- <i>x</i> /0/29: <i>[0-2]</i>	et- <i>x</i> /0/29
	et-x/0/30: [0-7]	et- <i>x</i> /0/30: <i>[0-7]</i>	et- <i>x</i> /0/30	et- <i>x</i> /0/30: <i>[0-7]</i>	et-x/0/30 et-x/0/30:[0-1] et-x/0/30:[0-3]	et- <i>x</i> /0/30: <i>[0-2]</i>	et- <i>x</i> /0/30
	et-x/0/31: [0-7]	et-x/0/31: [0-7]	et- <i>x</i> /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- <i>x</i> /0/32: <i>[0-7]</i>	et- <i>x</i> /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- <i>x</i> /0/32
	et-x/0/33: [0-7]	et- <i>x</i> /0/33: <i>[0-7]</i>	et- <i>x</i> /0/33	et- <i>x</i> /0/33: <i>[0-7]</i>	et-x/0/33 et-x/0/33:[0-1] et-x/0/33:[0-3]	et- <i>x</i> /0/33: <i>[0-2]</i>	et- <i>x</i> /0/33
	et-x/0/34: [0-7]	et- <i>x</i> /0/34: <i>[0-7]</i>	et- <i>x</i> /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- <i>x</i> /0/34

Table 118: 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- <i>x</i> /0/35: <i>[0-7]</i>	et- <i>x</i> /0/35: <i>[0-7]</i>	et- <i>x</i> /0/35	et- <i>x</i> /0/35: <i>[0-7]</i>	et-x/0/35 et-x/0/35:[0-1] et-x/0/35:[0-3]	et- <i>x</i> /0/35: <i>[0-2]</i>	et- <i>x</i> /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 291. 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 292.

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 119 on page 383 provides the basic details of PTX10001-36MR router.

Table 119: 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 119: PTX10001-36MR Router Details and Description (Continued)

Details	Description
MTU size	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.
Forward Error Correction (FEC) support	 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. For information about FEC support, see fec (gigether).

Table 120: 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 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	400-Gbps
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 121: Channelization and FEC support on 400- and 200-Gigabit Ethernet capable ports

Speed Supported	You can channelize to:
400-Gigabit Ethernet capable ports	 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	 You can configure all four 100-Gigabit Ethernet capable ports (ports 4, 5, 6, and 7) to operate at 100-Gbps speeds. 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: 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 289 in 20.2 releases. See *speed* for more details.

From Junos OS Evolved Release 20.3R1 you must use, set interfaces $et - \frac{fpc}{\sqrt{pic}}$ speed <10g | 25g | 50g | 40g | 100g | 200g | 400g>, or set interfaces $et - \frac{fpc}{\sqrt{pic}}$ 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 124 on page 389 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 122 on page 387 provides you the ports that you must power-off while configuring different speeds.

Table 122: 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 0/1/4 and 0/1/6 0/2/4 and 0/2/6	0/0/5 and 0/0/7 0/1/5 and 0/1/7 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 123 on page 387:

Table 123: Number of sub-ports supported for a particular speed

Port speed	Valid values for number-of-sub-ports
200-Gbps	1, 2

Table 123: Number of sub-ports supported for a particular speed (Continued)

Port speed	Valid values for number-of-sub-ports
100-Gbps QSFP56-DD-400G-ZR-M and QSFP56-DD-400G-ZR-M-HP support 3x100G channelisation.	1, 2, 4 3
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 124: Interface Naming Convention for PTX10001-36MR Router

PIC	4x100- Gigabit Ethernet Interface	4x10- Gigabit Ethernet Interface 4x25- Gigabit Ethernet Interface	2x100- Gigabit Ethernet Interface 2x50- Gigabit Ethernet Interface 2x200- Gigabit Ethernet Interface	8x25- Gigabit Ethernet Interface 8x50- Gigabit Ethernet Interface	1x400- Gigabit Ethernet Interface 1x200- 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
	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
	Unsupport ed	et-0/0/4: [0-3]	Unsupport ed	Unsupport ed	Unsupport ed	et-0/0/4	et-0/0/4
	Unsupport ed	Unused	Unsupport ed	Unsupport ed	Unsupport ed	et-0/0/5	Unused
	Unsupport ed	et-0/0/6: [0-3]	Unsupport ed	Unsupport ed	Unsupport ed	et-0/0/6	et-0/0/6
	Unsupport ed	Unused	Unsupport ed	Unsupport ed	Unsupport ed	et-0/0/7	Unused

Table 124: Interface Naming Convention for PTX10001-36MR Router (Continued)

PIC	4x100- Gigabit Ethernet Interface	4x10- Gigabit Ethernet Interface 4x25- Gigabit Ethernet Interface	2x100- Gigabit Ethernet Interface 2x50- Gigabit Ethernet Interface 2x200- Gigabit Ethernet Interface	8x25- Gigabit Ethernet Interface 8x50- Gigabit Ethernet Interface	1x400- Gigabit Ethernet Interface 1x200- Gigabit Ethernet Interface	1x100- Gigabit Ethernet Interface	1x40- Gigabit Ethernet Interface
	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
	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

Table 124: Interface Naming Convention for PTX10001-36MR Router (Continued)

PIC	4x100- Gigabit Ethernet Interface	4x10- Gigabit Ethernet Interface 4x25- Gigabit Ethernet Interface	2x100- Gigabit Ethernet Interface 2x50- Gigabit Ethernet Interface 2x200- Gigabit Ethernet Interface	8x25- Gigabit Ethernet Interface 8x50- Gigabit Ethernet Interface	1x400- Gigabit Ethernet Interface 1x200- Gigabit Ethernet Interface	1x100- Gigabit Ethernet Interface	1x40- Gigabit Ethernet Interface
	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
	Unsupport ed	et-0/1/4: [0-3]	Unsupport ed	Unsupport ed	Unsupport ed	et-0/1/4	et-0/1/4
	Unsupport ed	Unused	Unsupport ed	Unsupport ed	Unsupport ed	et-0/1/5	Unused
	Unsupport ed	et-0/1/6: [0-3]	Unsupport ed	Unsupport ed	Unsupport ed	et-0/1/6	et-0/1/6
	Unsupport ed	Unused	Unsupport ed	Unsupport ed	Unsupport ed	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

Table 124: Interface Naming Convention for PTX10001-36MR Router (Continued)

PIC	4x100- Gigabit Ethernet Interface	4x10- Gigabit Ethernet Interface 4x25- Gigabit Ethernet Interface	2x100- Gigabit Ethernet Interface 2x50- Gigabit Ethernet Interface 2x200- Gigabit Ethernet Interface	8x25- Gigabit Ethernet Interface 8x50- Gigabit Ethernet Interface	1x400- Gigabit Ethernet Interface 1x200- Gigabit Ethernet Interface	1x100- Gigabit Ethernet Interface	1x40- Gigabit Ethernet Interface
	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
	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
	Unsupport ed	et-0/2/4: [0-3]	Unsupport ed	Unsupport ed	Unsupport ed	et-0/2/4	et-0/2/4
	Unsupport ed	Unused	Unsupport ed	Unsupport ed	Unsupport ed	et-0/2/5	Unused
	Unsupport ed	et-0/2/6: [0-3]	Unsupport ed	Unsupport ed	Unsupport ed	et-0/2/6	et-0/2/6

Table 124: Interface Naming Convention for PTX10001-36MR Router (Continued)

PIC	4x100- Gigabit Ethernet Interface	4x10- Gigabit Ethernet Interface 4x25- Gigabit Ethernet Interface	2x100- Gigabit Ethernet Interface 2x50- Gigabit Ethernet Interface 2x200- Gigabit Ethernet Interface	8x25- Gigabit Ethernet Interface 8x50- Gigabit Ethernet Interface	1x400- Gigabit Ethernet Interface 1x200- Gigabit Ethernet Interface	1x100- Gigabit Ethernet Interface	1x40- Gigabit Ethernet Interface
	Unsupport ed	Unused	Unsupport ed	Unsupport ed	Unsupport ed	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
	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 291. 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 292.

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 125: 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 Channelized Interfaces" on page 292 and "Configure Port Speed for Non-Channelized Interfaces" on page 291 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 126 on page 395:

Table 126: 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 127 on page 396 lists the naming conventions used for interfaces on PTX10K-LC1202-36MR (for 100-Gbps, 2x100-Gbps, 400-Gpbs, and 4x100-Gbps speeds) for PTX10008 routers. PTX10008 routers support eight PTX10K-LC1202-36MR line cards.

Table 127: 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- <i>x</i> /0/0	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/1	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/2	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/3	Not Supported	Not Supported	Not Supported
	et-x/0/4	et- <i>x</i> /0/4: <i>[0-1]</i>	et- <i>x</i> /0/4	et- <i>x</i> /0/4: <i>[0-3]</i>
	et- <i>x</i> /0/5	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/6	Not Supported	Not Supported	Not Supported

Table 127: 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- <i>x</i> /0/7	Not Supported	Not Supported	Not Supported
	et-x/0/8	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/9	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/10	et- <i>x</i> /0/10: <i>[0-1]</i>	et- <i>x</i> /0/10	et- <i>x</i> /0/10: <i>[0-3]</i>
	et- <i>x</i> /0/11	Not Supported	Not Supported	Not Supported
	et-x/0/12	Not Supported	Not Supported	Not Supported
	et-x/0/13	Not Supported	Not Supported	Not Supported
	et-x/0/14	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/15	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/16	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/17	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/18	Not Supported	Not Supported	Not Supported
	et-x/0/19	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/20	Not Supported	Not Supported	Not Supported

Table 127: 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- <i>x</i> /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- <i>x</i> /0/24: <i>[0-1]</i>	et- <i>x</i> /0/24	et- <i>x</i> /0/24: <i>[0-3]</i>
	et-x/0/25	Not Supported	Not Supported	Not Supported
	et-x/0/26	Not Supported	Not Supported	Not Supported
	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- <i>x</i> /0/30: <i>[0-1]</i>	et- <i>x</i> /0/30	et- <i>x</i> /0/30: <i>[0-3]</i>
	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

Table 127: Interface Naming Convention for PTX10K-LC1202-36MR line card for PTX10008 Routers - Speeds 100-Gbps, 2x100-Gbps, 400-Gpbs, and 4x100-Gbps (*Continued*)

PIC	100-Gigabit Ethernet	2x100-Gigabit	400-Gigabit Ethernet	4x100-Gigabit
	Interface	Ethernet Interface	Interface	Ethernet Interface
	et- <i>x</i> /0/35	Not Supported	Not Supported	Not Supported

Table 128 on page 399 lists the naming conventions used for interfaces on PTX10K-LC1202-36MR (for 10-Gbps, 4x10-Gbps, 4v25-Gbps, and 8x25-Gbps speeds) for PTX10008 routers. PTX10008 routers support eight PTX10K-LC1202-36MR line cards.

Table 128: 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- <i>x</i> /0/0	et- <i>x</i> /0/0: <i>[0-3]</i>	et- <i>x</i> /0/0	et- <i>x</i> /0/0: <i>[0-3]</i>	Not Supported
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et-x/0/2	et- <i>x</i> /0/2: <i>[0-3]</i>	et- <i>x</i> /0/2	et- <i>x</i> /0/2: <i>[0-3]</i>	Not Supported
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et-x/0/4	et- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4	et- <i>x</i> /0/4: <i>[0-3]</i>	et- <i>x</i> /0/4: <i>[0-7]</i>
	et- <i>x</i> /0/5	et- <i>x</i> /0/5: <i>[0-3]</i>	et- <i>x</i> /0/5	et- <i>x</i> /0/5: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/6	et- <i>x</i> /0/6: <i>[0-3]</i>	et- <i>x</i> /0/6	et- <i>x</i> /0/6: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/7	et- <i>x</i> /0/7: <i>[0-3]</i>	et- <i>x</i> /0/7	et- <i>x</i> /0/7: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/8	et- <i>x</i> /0/8: <i>[0-3]</i>	et- <i>x</i> /0/8	et- <i>x</i> /0/8: <i>[0-3]</i>	Not Supported

Table 128: 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- <i>x</i> /0/9	et-x/0/9:[0-3]	et- <i>x</i> /0/9	et-x/0/9:[0-3]	Not Supported
	et- <i>x</i> /0/10	et- <i>x</i> /0/10: <i>[0-3]</i>	et- <i>x</i> /0/10	et- <i>x</i> /0/10: <i>[0-3]</i>	et- <i>x</i> /0/10: <i>[0-7]</i>
	et- <i>x</i> /0/11	et- <i>x</i> /0/11: <i>[0-3]</i>	et- <i>x</i> /0/11	et- <i>x</i> /0/11: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/12	et- <i>x</i> /0/12: <i>[0-3]</i>	et- <i>x</i> /0/12	et- <i>x</i> /0/12: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/13	et- <i>x</i> /0/13: <i>[0-3]</i>	et- <i>x</i> /0/13	et- <i>x</i> /0/13: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/14	et- <i>x</i> /0/14: <i>[0-3]</i>	et- <i>x</i> /0/14	et- <i>x</i> /0/14: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/15	et- <i>x</i> /0/15: <i>[0-3]</i>	et- <i>x</i> /0/15	et- <i>x</i> /0/15: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/16	et- <i>x</i> /0/16: <i>[0-3]</i>	et- <i>x</i> /0/16	et- <i>x</i> /0/16: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/17	et- <i>x</i> /0/17: <i>[0-3]</i>	et- <i>x</i> /0/17	et- <i>x</i> /0/17: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/18	et- <i>x</i> /0/18: <i>[0-3]</i>	et- <i>x</i> /0/18	et- <i>x</i> /0/18: <i>[0-3]</i>	Not Supported
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/20	et- <i>x</i> /0/20: <i>[0-3]</i>	et- <i>x</i> /0/20	et- <i>x</i> /0/20: <i>[0-3]</i>	Not Supported
	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	et- <i>x</i> /0/22	et- <i>x</i> /0/22: <i>[0-3]</i>	et- <i>x</i> /0/22	et- <i>x</i> /0/22: <i>[0-3]</i>	Not Supported

Table 128: 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/0/23	et- <i>x</i> /0/23: <i>[0-3]</i>	et-x/0/23	et- <i>x</i> /0/23: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/24	et- <i>x</i> /0/24: <i>[0-3]</i>	et- <i>x</i> /0/24	et- <i>x</i> /0/24: <i>[0-3]</i>	et- <i>x</i> /0/24: <i>[0-7]</i>
	et-x/0/25	et- <i>x</i> /0/25: <i>[0-3]</i>	et- <i>x</i> /0/25	et- <i>x</i> /0/25: <i>[0-3]</i>	Not Supported
	et-x/0/26	et- <i>x</i> /0/26: <i>[0-3]</i>	et- <i>x</i> /0/26	et- <i>x</i> /0/26: <i>[0-3]</i>	Not Supported
	et-x/0/27	et- <i>x</i> /0/27: <i>[0-3]</i>	et- <i>x</i> /0/27	et- <i>x</i> /0/27: <i>[0-3]</i>	Not Supported
	et-x/0/28	et- <i>x</i> /0/28: <i>[0-3]</i>	et- <i>x</i> /0/28	et- <i>x</i> /0/28: <i>[0-3]</i>	Not Supported
	et-x/0/29	et- <i>x</i> /0/29: <i>[0-3]</i>	et- <i>x</i> /0/29	et- <i>x</i> /0/29: <i>[0-3]</i>	Not Supported
	et-x/0/30	et- <i>x</i> /0/30: <i>[0-3]</i>	et- <i>x</i> /0/30	et- <i>x</i> /0/30: <i>[0-3]</i>	et- <i>x</i> /0/30: <i>[0-7]</i>
	et- <i>x</i> /0/31	et- <i>x</i> /0/31: <i>[0-3]</i>	et- <i>x</i> /0/31	et- <i>x</i> /0/31: <i>[0-3]</i>	Not Supported
	et-x/0/32	et- <i>x</i> /0/32: <i>[0-3]</i>	et- <i>x</i> /0/32	et- <i>x</i> /0/32: <i>[0-3]</i>	Not Supported
	et-x/0/33	et- <i>x</i> /0/33: <i>[0-3]</i>	et- <i>x</i> /0/33	et- <i>x</i> /0/33: <i>[0-3]</i>	Not Supported
	et-x/0/34	et- <i>x</i> /0/34: <i>[0-3]</i>	et- <i>x</i> /0/34	et- <i>x</i> /0/34: <i>[0-3]</i>	Not Supported
	et- <i>x</i> /0/35	et- <i>x</i> /0/35: <i>[0-3]</i>	et- <i>x</i> /0/35	et- <i>x</i> /0/35: <i>[0-3]</i>	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 then 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 129 on page 402 summarizes the Packet Forwarding Engine mapping and the supported port speeds for PTX10003-80C.

Table 129: Port Speed for PTX10003-80C

FPC	PIC	Ports	Port Speed Supported
FPC0	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

Table 129: Port Speed for PTX10003-80C (Continued)

FPC	PIC	Ports	Port Speed Supported
FPC1	PICO	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

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 130: Port Speed for PTX10003-160C

FPC	PIC	Ports	Port Speed Supported
FPC0	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

Table 130: Port Speed for PTX10003-160C (Continued)

FPC	PIC	Ports	Port Speed Supported
FPC1	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
FPC2	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
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 131: 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 132: 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
Port 9	400 Gbps

Table 133: 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

Table 133: Configuration Guidelines for 4x100 Gbps (PTX10003) (Continued)

Ports with speed of 4x100 Gbps	Ports configured as Unused
Port 9	Port 7 and Port 8

Table 134: 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
	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

Table 134: Interface Naming Convention for PTX10003 (Continued)

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
	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 288 . For information on how to configure the speed at the port level, see "Configure Port Speed - Port Level" on page 289 .



400G ZR and 400G ZR+ Optical Transceivers

Overview | 410

400ZR and 400G OpenZR+ | 412

Configure 400ZR and 400G OpenZR+ Coherent Optics | 423

Overview

IN THIS SECTION

Benefits of 400ZR and 400G OpenZR+ Optics | 411

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.

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.

• **OdBm output transmit power**: QDD-400G-ZR-M-HP models can operate in high power mode and deliver a transmit power output of 0 dBm.

Table 135: Optics-Supported Devices

Feature	Supported Devices	Initial Release
"Application Selection" on page 413	PTX10001-36MR	Junos OS Evolved 23.1 R1
"Target Output Power" on page 418	PTX10001-36MR	Junos OS Evolved 23.1 R1

Table 135: Optics-Supported Devices (Continued)

Feature	Supported Devices	Initial Release
"Enhanced Loopback Options" on page 420	PTX10001-36MR	Junos OS Evolved 23.1 R1

400ZR and 400G OpenZR+

IN THIS SECTION

- Licensing | 412
- Application Selection | 413
- Configure Application Selection | 415
- Target Output Power | 418
- Configure Target Output Power | 418
- Enhanced Loopback Options | 420
- Tunable Wavelength | 421
- Signal Power Performance | 421

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.

Application Selection

The 400ZRand 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 for devices that support application selection.

Table 136: 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
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 137: 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

Table 137: Supported Applications of 400G OpenZR+ Optics (Continued)

Application	Host ID	Host Interface Code Description	Media ID	Media Interface Code Description	Channelization
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
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 138: Default Application for 400ZR and 400G OpenZR+

Optics	Speed	Default Application
400GbE ZR	1x400G	CFEC amplified
400GbE ZR+	1x400G	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- \Leftrightarrow optics-options tx-power \Leftrightarrow command. Changing or adjusting to a different Tx power improves transmission quality. See tx-power.

Table 139: Default Output Power

Transceiver	Default Transmit Power	
QDD-400G-ZR	-10 dBm	
JCO400-QDD-ZR	-10 dBm	
QDD-400G-ZR-M	-10 dBm	
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
```

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 the full C-band International Telecommunication Union (ITU)-Grid tunable optics.

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- \Leftrightarrow 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 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 : Off Module temperature high alarm : Off Module temperature low alarm : Off Module temperature high warning Module temperature low warning : Off Module voltage high alarm 0ff

```
: Off
 Module voltage low alarm
 Module voltage high warning
                                              0ff
                                           : Off
 Module voltage low warning
 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
                                           : 0.006 mW / -22.00 dBm
 Laser rx power low alarm threshold
 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
                                           : 48 degrees C / 118 degrees F
 Laser temperature
                                              0.887 mW / -0.51 dBm
 Laser receiver power
 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
                                           : Off
 Laser bias current low alarm
  Laser bias current high warning
                                           : Off
```

```
Laser bias current low warning
                              : Off
                                    : Off
Laser temperature high alarm
Laser temperature low alarm
                                    : Off
                               : Off
Laser temperature high warning
Laser temperature low warning
                                  : Off
Laser receiver power high alarm
                                  : Off
                                   : Off
Laser receiver power low alarm
Laser receiver power high warning : Off
Laser receiver power low warning
                                  : Off
                           : Off
: Off
: Off
Laser output power high alarm
Laser output power low alarm
Laser output power high warning
Laser output power low warning
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm
                                    : Off
                                    : Off
Tx laser disabled alarm
```

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



100GbE, 40GbE, and 10GbE Optical Transceivers

100GbE, 40GbE, and 10GbE | 426

100GbE, 40GbE, and 10GbE

IN THIS SECTION

- Overview | 426
- Supported Features | 428
- Supported PICs | 432
- Configure Optics | 433
- Supported OTN and Optics Options | 438
- Forward Error Correction and Bit Error Rate | 475
- | 486
- | 487
- | 487
- | 487
- | 487
- Interface Mapping and Modulation format for PTX10K-LC1104 Line Card | 487

Read this topic for information about optics support that the specific line cards and devices provide.

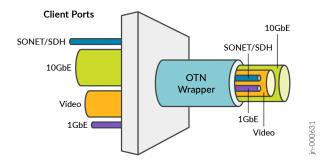
Overview

IN THIS SECTION

Benefits | 427

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 particular 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 the following:

- 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

- Allows FEC to improve the system performance.
- Provides alarm handling capability.
- Protects the network from adverse "mixing" of heterogeneous services, saves bandwidth, and transfers any traffic transparently.

• Easily scalable, supports dedicated Ethernet links at 1GbE, 10GbE, 40GbE, and 100GbE.

Supported Features

Table 140: Supported optics features

Features	Description
Interface	 Supports: 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.
	 Attributes: 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: Set functionality Juniper Networks Black-Link MIB Optics MIB FRU MIB

Table 140: Supported optics features (Continued)

Features	Description
	Performance monitoring (15-minute and 1-day performance monitoring and historic statistics): Near-end and far-end TCA BER FEC Optical
L2 and L3 Features	 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	You receive TCA when you cross a certain configurable threshold (near-end measurement threshold or far-end measurement threshold). The TCA remains so until the end of the 15-minute interval for parameters such as OTU and ODU

Table 140: Supported optics features (Continued)

Features	Description
	 Supported alarms: Background block error threshold (BBE) Errored seconds (ES) threshold Severely errored second (SES) threshold Unavailable seconds (UAS) threshold

OTN Alarms and Defects

Table 141: 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

Table 141: OTN Alarms and Defects (Continued)

OTN Alarms and Defects	Description
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]
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

Table 141: OTN Alarms and Defects (Continued)

OTN Alarms and Defects	Description
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 142: 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

Table 142: Supported PICs (Continued)

PIC	
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 |
odu-iae | odu-lck | odu-oci | odu-sd | odu-ttim |opu-ptim | otu-ais | otu-bdi | otu-fec-deg
| otu-fec-exe | otu-iae | 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 stay 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 stay 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 432		
optics-opti	ons		
otn-options			
signal-degr	nde		
preemptive	-fast-reroute		
wavelength			

Supported OTN and Optics Options

IN THIS SECTION

- Supported OTN Options on ACX6360 and ACX5448 Routers | 438
- Supported OTN Options on MX Series Routers | 444
- Supported OTN Options on PTX Series Routers | 454
- Supported Optics Options on ACX6360 and ACX5448-D Routers | 465
- Supported Optics Options on PTX10008 and PTX10016 Series Routers | 469

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 143 on page 438 lists the statements that are supported on 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 143: 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

Table 143: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
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
	signal- degrade- monitor- enable no-signal- degrade- monitor- enable	Yes	Yes
	odu-backward- frr-enable no - odu-backward- frr-enable	No	No

Table 143: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
	odu-signal-degrade- monitor- enable no -odu-signal-degrade- monitor-enable	No	No
tca	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
	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

Table 143: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Statement Options		ACX5448 (19.2R1)
	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
	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

Table 143: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
trigger <i>trigger-</i> <i>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
	odu-oci (hold-time (down up) ignore)	Yes	Yes
	odu-sd (hold-time (down up) ignore)	Yes	Yes
	odu-ttim	Yes	Yes

Table 143: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
	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
	otu-iae (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 identifier	Yes	Yes
	odu-expected- receive-sapi identifier	Yes	Yes
	odu-sapi <i>identifier</i>	Yes	Yes
	otu-dapi <i>identifier</i>	Yes	Yes

Table 143: Statements Supported on ACX6360 and ACX5448 Routers (Continued)

Statement	Options	ACX6360 (18.3R1)	ACX5448 (19.2R1)
	otu-expected- receive-dapi identifier	Yes	Yes
	otu-expected- receive-sapi identifier	Yes	Yes
	otu-sapi <i>identifier</i>	Yes	Yes

SEE ALSO

Configuring OTN

Supported OTN Options on MX Series Routers

Table 144 on page 444 lists the statements that are supported on 100-Gigabit Ethernet MICs on MX Series routers at the [edit interfaces *interface-Nome* otn-options] hierarchy level.

Table 144: 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

Table 144: 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)
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
	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

Table 144: 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)
	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
	odu-backward-frr-eNoble no-odu-backward-frr-eNoble	No	Yes
	odu-sigNol-degrade- monitor-eNoble no -odu- sigNol-degrade- monitor- eNoble	No	Yes

Table 144: 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)
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
	q-threshold-sigNol-degrade- clear	No	Yes
tca	odu-tca-bbe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes

Table 144: 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-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
	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

Table 144: 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-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
	otu-tca-uas-fe (eNoble-tca no-eNoble-tca threshold threshold-24hrs)	Yes	Yes
transport- monitoring	-	No	Yes

Table 144: 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)
trigger <i>trigger-</i> <i>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
	odu-iae (hold-time (down up) ignore)	Yes	Yes
	odu-lck (hold-time (down up) ignore)	Yes	Yes

Table 144: 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-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
	otu-fec-deg (hold-time (down up) ignore)	No	Yes
	otu-fec-exe (hold-time (down up) ignore)	No	Yes

Table 144: 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-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>identifier</i> y	Yes	Yes
	odu-dapi-first-byte-nul no-odu-dapi-first-byte-nul	No	Yes
	odu-expected-receive- dapi <i>identifier</i>	Yes	Yes
	odu-expected-receive-dapi- first-byte-nul no-odu- expected-receive-dapi- first-byte-nul	No	Yes

Table 144: 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-sapi identifier	Yes	Yes
	odu-sapi <i>identifier</i> y	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 identifier	Yes	Yes
	otu-expected-receive-dapi- first-byte-nul no-odu- expected-receive-dapi- first-byte-nul	No	Yes
	otu-expected- receive-sapi identifier	Yes	Yes
	otu-expected-receive-sapi- first-byte-nul no-odu- expected-receive-dapi- first-byte-nul	No	Yes

Table 144: 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-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 145 on page 454 lists the statements that are supported on 100-Gigabit Ethernet PICs on PTX Series routers at the [edit interfaces *interface-name* otn-options] hierarchy level.

Table 145: 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-24-1 OG-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
fec	(efec gfec gfec-sdfec none ufec)	Yes	Yes (gfec, none)	Yes	Yes

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
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
	monitor- end-point no-monitor-end- point	Yes	No	No	Yes

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	number- of-frames value	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 value	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
otu-ttim- action- enable no- otu-ttim- action- enable	-	Yes	Yes	Yes	Yes

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
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
	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

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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 value	Yes	Yes	Yes	Yes
	interval <i>value</i>	Yes	Yes	Yes	Yes
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

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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
	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

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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
	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

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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
	otu-tca-uas-fe (enable-tca no- enable-tca threshold threshold-24hrs)	Yes	No	Yes	Yes
transport- monitoring	-	No	No	Yes	Yes
trigger trigger- identifier	oc-lof (hold-time (down up) ignore)	Yes	Yes	Yes	Yes

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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
	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

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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
	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

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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
	otu-ttim (hold- time (down up) ignore)	Yes	Yes	Yes	Yes
tti <i>tti-</i> identifier	odu-dapi <i>identifier</i>	Yes	14.1R2 14.2	Yes	Yes
	odu-expected- receive-dapi <i>identifier</i>	Yes	Yes	Yes	Yes

Table 145: 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 OG-W- SFPP (PTX5000) (14.2R1)	PTX10K- LC1104 (PTX10008 and PTX10016) (18.3R1)
	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
	otu-sapi <i>identifier</i>	Yes	Yes	Yes	Yes

Configuring OTN Interfaces on P1-PTX-2-100G-WDM

Supported Optics Options on ACX6360 and ACX5448-D Routers

Table 146 on page 466 lists the statements that are supported on ACX6360 and ACX5448-D routers at the [edit interfaces *interface-name* optics-options] hierarchy level.

Table 146: Statements Supported on ACX6360 and ACX5448-D Routers

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1- S1)	Interfaces Supported
fec	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
modulation-format	(16qam 8qam qpsk)	Yes	ot
signal-degrade	ber-threshold-clear <i>value</i>	Yes	ot
	ber-threshold-signal-degrade value	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

Table 146: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1- S1)	Interfaces Supported
	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	
	fec-corrected-errors-high-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
	fec-ucorrected-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	

Table 146: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1- S1)	Interfaces Supported
	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	
	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	
	<pre>snr-low-tca (enable-tca no- enable-tca threshold threshold-24hrs)</pre>	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	

Table 146: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

Statement	Options	Release (18.2R1, 18.3R1, and 19.2R1- S1)	Interfaces Supported
	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	
	tx-power-low-tca (enable-tca no-enable-tca threshold threshold-24hrs)	Yes	
tx-power	dbm	Yes	ot
wavelength	nm	Yes	ot

Configuring OTN

Supported Optics Options on PTX10008 and PTX10016 Series Routers

Table 147 on page 470 lists the statements that are supported on PTX10008 and PTX10016 Series routers at the [edit interfaces *interface-name* optics-options] hierarchy level.

Table 147: 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	<pre>(enable-tca no-enable-tca threshold threshold-24hrs)</pre>	Yes	ot	
tx-power	dbm	Yes	ot	
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 value	Yes		
	interval <i>value</i>	Yes		
tca	odu-tca-bbe (enable-tca no- enable-tca threshold threshold-24hrs)	Yes	ot	

Table 147: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

Statement	Options	Release (18.3R1)	Interfaces Supported
	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	
	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	

Table 147: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

Statement	Options	Release (18.3R1)	Interfaces Supported
	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	
	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-</i> identifier	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	

Table 147: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

Table 147. Statement	s Supported on PTA10008 an	la PIX10016 Series Routers (C	Ontinueu)
Statement	Options	Release (18.3R1)	Interfaces Supported
	oc-tsf (hold-time (down up) ignore)	Yes	
	oc-wavelength-lock (hold-time (down up) ignore)	Yes	
	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	

Table 147: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

Statement	Options	Release (18.3R1)	Interfaces Supported
	otu-ais (hold-time (down up) ignore)	Yes	
	otu-bdi (hold-time (down up) ignore)	Yes	
	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 | 475
- Supported Forward Error Correction Modes | 480
- ODU Path Delay Measurement for Performance Monitoring | 481

OTN interfaces use pre-forward error correction (pre-FEC) BER for monitoring the condition of an OTN link. Use 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 | 478

Optical interfaces on PTX Series Packet Transport Routers 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:

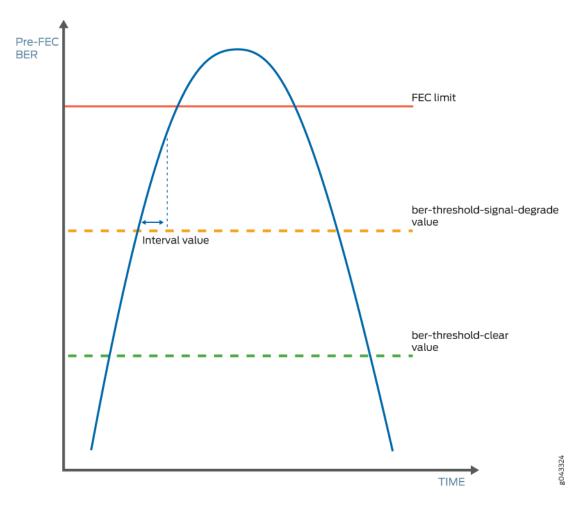
- P1-PTX-2-100G-WDM
- P2-100GE-OTN
- P1-PTX-24-10G-W-SFPP

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 this pre-FEC BER threshold logic is combined with MPLS fast reroute, then packet loss can be minimized or prevented.

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 476. After an alarm is raised, if the BER returns to a level below the threshold clear value (*ber-threshold-clear*), the alarm is cleared.





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 trasmitted 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 Q^2 -factor measurements (also known as Q-factor). Q^2 -factor is expressed in units of decibels relative to a Q^2 -factor of zero (dBQ). Q^2 -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 Q^2 -factor to BER to enter them in the CLI by using scientific notation. BER can be converted to Q^2 -factor by using the following equation:

```
Q^2-factor = 20 * log10 (sqrt(2) * erfcinv(2 * BER))
```

TIP: To convert between Q^2 -factor and BER in a spreadsheet program, you can approximate the values by using the following formulas:

• To calculate Q²-factor:

```
= 20 * LOG10(-NORMSINV(BER))
```

• To calculate BER:

= 1 - NORMSDIST($10^{0.05} * Q2$ -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 148 on page 478 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 148: 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

Table 148: Example—Signal Degrade and Clear Threshold Values at 1 dBQ (Continued)

PIC	FEC Type	FEC Limit		Signal Degrade Threshold		Clear Threshold	
		Q ² -Factor	BER	Q ² -Factor	BER	Q ² -Factor	BER
	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 149 on page 479 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 149: 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 150: 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)
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 151: 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 152: 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
, ACX5448-D	sdfec15	8-QAM	200G
, ACX5448-D	sdfec15	16-QAM	200G

fec

ODU Path Delay Measurement for Performance Monitoring

IN THIS SECTION

- Overview | 482
- Enabling ODU Path Delay Measurement | 483
- Disabling ODU Path Delay Measurement | 485

Read this topic to understand ODU path delay measurement and performance monitoring.

Overview

IN THIS SECTION

Guidelines to Configure Delay Measurement | 482

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.

```
Frame# 10 9 8 7 6 5 4 3 2 1
Tx DM bit 1 1 1 1 1 1 1 0 0
Rx DM bit 1 1 1 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 inservice 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 re-enable
 measurement of delay 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 | 426

Supported OTN and Optics Options

RELATED DOCUMENTATION

100GbE, 40GbE, and 10GbE | 426

Configuring OTN

SEE ALSO

Configuring OTN

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
	16QAM	et-x/0/0 et-x/0/1

(Continued)

"ot-" interface	Modulation Format	Mapped "et" interface(s)
ot-0/0/1	QPSK	et-x/0/2
	8QAM	et-x/0/1 et-x/0/2
		et-x/0/2
	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
	16QAM	et-x/0/4
		et-x/0/5
ot-0/0/3	QPSK	et-x/0/6
	8QAM	et-x/0/5
		et-x/0/6
	16QAM	et-x/0/6
		et-x/0/7
ot-0/0/4	QPSK	et-x/0/8
	8QAM	et-x/0/8
		et-x/0/9

(Continued)

"ot-" interface	Modulation Format	Mapped "et" interface(s)
	16QAM	et-x/0/8 et-x/0/9
ot-0/0/5	QPSK	et-x/0/10
	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.

RELATED DOCUMENTATION

Configuring OTN

Forward Error Correction (FEC) and Bit Error Rate (BER)



Troubleshooting Information

Monitoring and Troubleshooting Ethernet Interfaces | 492

Monitoring and Troubleshooting Ethernet Interfaces

IN THIS CHAPTER

- Passive Monitoring on Ethernet Interfaces Overview | 492
- Enabling Passive Monitoring on Ethernet Interfaces | 494
- Link Degrade Monitoring Overview | 497
- Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501
- Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 517
- Performing Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 538
- Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 558
- Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters | 566
- Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State | 573
- Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | 575

Passive Monitoring on Ethernet Interfaces Overview

The Monitoring Services I and Monitoring Services II PICs are designed to enable IP services. You can monitor IPv4 traffic if you have a Monitoring Services PIC installed in the router with the following PICs:

- 10-port Gigabit Ethernet PIC with SFPs
- 4-port Gigabit Ethernet PIC with SFPs
- · 2-port Gigabit Ethernet PIC with SFPs
- 1-port 10-Gigabit Ethernet PIC

NOTE: The PICs in the preceding list support only IPv4.

NOTE: Starting with Junos OS Release 9.5, I2.0 based M120 routers and I3.0 based M320 routers with the PICs in the preceding list support passive monitoring. Other M Series and T Series routers with the PICs listed above started supporting passive monitoring before Junos OS Release 7.3. Support for 1-port 10-Gigabit Ethernet PIC with XENPAK on I2.0-based M120 routers and I3.0-based M320 routers was added in Junos OS Release 9.5.

 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (T640, T1600, and T4000 Core Routers) (supported on both WAN-PHY and LAN-PHY modes for both IPv4 and IPv6 addresses)

The following interfaces support passive monitoring on the I3.0-based MX 240, MX 480, and MX 960 routers, starting with Junos OS Release 8.5:

- Type 2 MX FPCs
- Type 3 MX FPCs
- Gigabit Ethernet Enhanced DPC with SFP (DPCE-R-40GE-SFP)
- 4-port 10-Gigabit Ethernet Enhanced DPCs with XFP (DPCE-R-4XGE-XFP)

The following interfaces support passive monitoring on the Trio-based MX 240, MX 480, and MX 960 routers:

- 10-Gigabit Ethernet MPC with SFP+
- 30-Gigabit Ethernet MPC
- 60-Gigabit Ethernet MPC

Passive monitoring is also supported on MX 80 routers with 10-Gigabit Ethernet MPC with SFP+ and 30-Gigabit Ethernet MPC interfaces.

Interfaces configured on the following FPCs and PIC support IPv6 passive monitoring on the T640, T1600, and T4000 routers:

- Enhanced Scaling FPC2
- Enhanced Scaling FPC3
- Enhanced Scaling FPC4
- Enhanced Scaling FPC4.1
- Enhanced II FPC1 (T640 and T1600 routers)
- Enhanced II FPC2 (T640 and T1600 routers)

- Enhanced II FPC3 (T640 and T1600 routers)
- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (supported on both WAN-PHY and LAN-PHY modes for both IPv4 and IPv6 addresses)
- Gigabit Ethernet PIC with SFP
- 10-Gigabit Ethernet PIC with XENPAK (T640 and T1600 routers)
- SONET/SDH OC192/STM64 PICs with XFP (T1600 and T4000 routers)
- SONET/SDH OC48c/STM16 PIC with SFP
- SONET/SDH OC12/STM4 (Multi-Rate) PIC with SFP (T1600 router)
- Type 1 SONET/SDH OC3/STM1 (Multi-Rate) PIC with SFP (T1600 router)

NOTE: Unlike IPv4 passive monitoring, IPv6 passive monitoring is not supported on Monitoring Services PICs. You must configure *port mirroring* to forward the packets from the passive monitored ports to other interfaces.

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
9.5	Starting with Junos OS Release 9.5, I2.0 based M120 routers and I3.0 based M320 routers with the PICs in the preceding list support passive monitoring.

RELATED DOCUMENTATION

Ethernet Interfaces User Guide for Routing Devices

Enabling Passive Monitoring on Ethernet Interfaces

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine silently drops packets coming from that interface and destined to the router itself. Passive monitoring mode also stops the Routing Engine from transmitting any packet from that interface. Packets received from the

monitored interface can be forwarded to monitoring interfaces. If you include the passive-monitor-mode statement in the configuration:

- Gigabit and Fast 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, go to the [edit interfaces interface-name] hierarchy level.

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

2. Include 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, go 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. Include 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 include 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, go 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. Include 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.

For information on FPCs and PICs that support IPv6 passive monitoring on the T640, T1600, and T4000 routers, see "Passive Monitoring on Ethernet Interfaces Overview" on page 492. Interfaces configured on these FPCs and PICs support IPv6 passive monitoring.

To configure port mirroring, include 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

Passive Monitoring on Ethernet Interfaces Overview | 492

Configuring Multiservice Physical Interface Properties

Junos OS Services Interfaces Library for Routing Devices

Ethernet Interfaces User Guide for Routing Devices

Link Degrade Monitoring Overview

IN THIS SECTION

Supported Platforms | 498

Link degrade monitoring enables you to monitor the quality of physical links on Ethernet interfaces and take corrective action when the link quality degrades beyond a certain level. You can configure this feature by applying the link-degrade-monitor statement at the [edit interfaces *interface-name*] hierarchy level.

NOTE: Ethernet links (using Link Fault Signaling (LFS) protocol), Layer 2 and Layer 3 protocols support physical link monitoring. However, these existing mechanisms cannot detect BER ranges that are very low (for example, 10⁻¹³ through 10⁻⁵).

When you configure link degrade monitoring on your device, this feature continuously monitors bit error rate (BER) value of the link and initiates a corrective action (media-based) when the BER value breaches a user-configured threshold. You can configure the threshold by using the thresholds statement at the [edit interfaces *interface-name* link-degrade-monitor] hierarchy level. The feature can detect a BER value as low as 10^{-13} through 10^{-5} , helping you prevent or minimize packet drops in physical links.

You can configure autorecovery or manual recovery method for the degraded link. In the case of manual recovery, you need to use the request interface link-degrade-recover *interface-name* statement to recover the degraded link. If autorecovery is configured, automatic recovery of the degraded link is attempted at the user configured intervals, and when the link's BER value is within the configured limit, the link is recovered.

When you configure link monitoring (by using the link-degrade-monitor statement), you can set the following options:

- BER threshold value at which the corrective action must be triggered on an interface. Use the thresholds (set | clear) statement to configure the BER threshold value or clear the BER threshold value.
- Link degrade interval value, which determines the number of consecutive link degradation events that are considered before any corrective action is taken. Use the thresholds interval *intervalvalue* statement to configure the interval.

- Warning threshold value for link degradation. Use the thresholds warning (set | clear) statement to configure the warning threshold value or clear the warning threshold value.
- Action that is taken when the configured BER threshold level is reached. Use the thresholds actions media-based statement to configure the action to be taken.
- Automatic or manual recovery. Use the recovery (auto | manual) statement to configure recovery.
- Time interval after which automatic recovery of the degraded link must be triggered. Use the recovery timer *timer-value* statement to configure the time interval.

Supported Platforms

Table 153 on page 499 and Table 154 on page 500 lists the platform series and line cards that support link degrade monitoring.

Table 153: MX Line Cards that Support Link Degrade Monitoring

Platform Series	MPC Line Cards Supported	DPC Line Cards Supported
MX	MPC4E-3D-2CGE-8XGE	DPCE-R-Q-4XGE-XFP
	MPC4E-3D-32XGE-SFPP	DPCE-R-4XGE-XFP
	MPC-3D-16XGE-SFP	DPCE-X-4XGE-XFP
	MPC3 with MIC3-3D-1X100GE-CFP	DPCE-X-Q-4XGE-XFP
	MPC3 with MIC3-3D-2X40GE-	DPCE-R-2XGE-XFP
	QSFPP	DPCE-R-4XGE-XFP
	MPC3 with MIC-3D-2XGE-XFP	DPCE-X-4XGE-XFP
	MPC3 with 2x10GE XFP MIC	On 10-Gigabit Ethernet interfaces:
	MPC3 with 2x10GE XFP MIC	DPCE-R-Q-20GE-2XGE
	MPC5 with following variants:	DPCE-R-20GE-2XGE
	• 2CGE + 4XGE	DPCE-X-20GE-2XGE
	• 24XGE + 6XLGE	
	MPC6 with the following variants:	
	• 2X100GE CFP2	
	• 24X10GE SFPP	
	24X10GE SFPP OTN	
	• 4x100GE CXP	
	MPC7E-MRATE	
	MPC7E-10G (non-MACsec mode)	
	MX2K-MPC8E with MIC- MRATE	

Table 153: MX Line Cards that Support Link Degrade Monitoring (Continued)

Platform Series	MPC Line Cards Supported	DPC Line Cards Supported
	 MX2K-MPC9E with MIC-MRATE NOTE: Link degrade monitoring is not supported on MACsecenabled MPC7E-10G and MIC-MACSEC-MRATE. MPC10E-10C-MRATE (10G, 25G, 40G, 100G and 400G interfaces) MPC10E-15C-MRATE (10G, 25G, 40G, 100G and 400G interfaces) MX2K-MPC11E (10G, 40G, 100G, and 400G interfaces) 	

Table 154: PTX Line Cards that Support Link Degrade Monitoring

Platform Series	Line Cards/PICs Supported
PTX3000 and PTX5000 (Junos)	P1-PTX-2-100GE-CFPP2-100GE-CFP2PF-1CGE-CFP
PTX10003 (Junos Evolved)	10G, 25G, 40G, 50G, 100G, 200G, and 400G interfaces
PTX10004 (Junos Evolved)	10G, 25G, 40G, 50G, 100G, 200G, and 400G interfaces
PTX10008 (Junos Evolved)	 JNP10K-LC1201 PTX10K-LC1202-36MR (10G and 100G interfaces)

Table 154: PTX Line Cards that Support Link Degrade Monitoring (Continued)

Platform Series	Line Cards/PICs Supported
PTX10001 (Junos Evolved)	Interface Naming Conventions for PTX10001-36MR Router

RELATED DOCUMENTATION

Physical Interface Damping Overview	
Fast Reroute Overview	
link-degrade-monitor	
thresholds	
recovery	
request interface link-degrade-recover	

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

IN THIS SECTION

- Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501
- Monitor Fast Ethernet and Gigabit Ethernet Interfaces | 503
- Fiber-Optic Ethernet Interface Specifications | 515

Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

IN THIS SECTION

Purpose | 502

Action | **502**

Meaning | 503

Purpose

To monitor Fast Ethernet and Gigabit Ethernet interfaces and begin the process of isolating interface problems when they occur.

Action

Table 155 on page 502 provides links and commands for monitoring Fast Ethernet and Gigabit Ethernet interfaces.

Table 155: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action
"Monitor Fast Ethernet and Gigabit Ethernet Interfaces" or	n page 503
"Display the Status of Fast Ethernet Interfaces" on page 501	show interfaces terse (fe* ge*)
1. "Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface" on page 503	show interfaces (fe-fpc/pic/port ge-fpc/pic/port)
"Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface" on page 503	show interfaces (fe-fpc/pic/port ge-fpc/pic/port) extensive
1. Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface	monitor interface (fe-fpc/pic/port ge-fpc/pic/port)
"Fiber-Optic Ethernet Interface Specifications" on page 515	

Meaning

You can use the above described commands to monitor and to display the configurations for Fast Ethernet and Gigabit Ethernet interfaces.

SEE ALSO

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 501

Monitor Fast Ethernet and Gigabit Ethernet Interfaces

IN THIS SECTION

- Display the Status of Fast Ethernet Interfaces | 503
- Display the Status of Gigabit Ethernet Interfaces | 505
- Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface | 507
- Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface | 509
- Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 513

By monitoring Fast Ethernet and Gigabit Ethernet interfaces, you begin to isolate Fast Ethernet and Gigabit Ethernet interface problems when they occur.

To monitor your Fast Ethernet and Gigabit Ethernet interfaces, follow these steps:

Display the Status of Fast Ethernet Interfaces

IN THIS SECTION

- Purpose | **504**
- Action | **504**
- Meaning | 504

Purpose

To display the status of Fast Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Action

```
user@host> show interfaces terse (fe* | ge*)
```

Sample Output

command-name

```
user@host> show interfaces terse fe*
Interface
              Admin Link Proto Local
                                                   Remote
fe-2/1/0
              up
                    up
fe-2/1/0.0
                    up inet 10.116.115.217/29
              up
fe-3/0/2
              up
                   down
fe-3/0/2.0
              up down
fe-3/0/3
              up
                    up
fe-3/0/3.0
                        inet 192.168.223.65/30
              up
                    up
fe-4/1/0
              down up
fe-4/1/0.0
                    down inet 10.150.59.133/30
              up
fe-4/1/1
                    up
              up
fe-4/1/1.0
              up
                    up
                       inet 10.150.59.129/30
                    down
fe-4/1/2
              up
fe-4/1/2.0
              up
                    down
```

Meaning

The sample output lists only the Fast Ethernet interfaces. It shows the status of both the physical and logical interfaces. For a description of what the output means, see Table 2.

Table 156: Status of Fast Ethernet Interfaces

Physical Interface	Logical Interface	Status Description
fe-2/1/0 Admin Up Link Up	fe-2/1/0.0 Admin Up Link Up	This interface has both the physical and logical links up and running.
fe-3/0/2 Admin Up Link Down	fe-3/0/2.0 Admin Up Link Down	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.
fe-4/1/0 Admin Down Link Up	fe-4/1/0.0 Admin Up Link Down	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.
fe-4/1/2 Admin Up Link Down	fe-4/1/2.0 Admin Up Link Down	This interface has both the physical and logical links down.

SEE ALSO

Monitoring Fast Ethernet and Gigabit Ethernet

Display the Status of Gigabit Ethernet Interfaces

IN THIS SECTION

• Purpose | **506**

Action | **506**

Meaning | 506

Purpose

To display the status of Gigabit Ethernet interfaces, use the following Junos OS command-line interface (CLI) operational mode command:

Action

Sample Output

user@host> show	interf	faces	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 157: Status of Gigabit Ethernet Interfaces

Physical Interface	Logical Interface	Status Description
ge-2/2/0 Admin Down Link Down	ge-2/2/0.0 Admin Up Link Down	This interface is administratively disabled (Admin Down). Both the physical and logical links are down (Link Down).
ge-2/3/0 Admin Up Link Up	ge-2/3/0.0 Admin Up Link Up	This interface has both the physical and logical links up and running.

Table 157: Status of Gigabit Ethernet Interfaces (Continued)

Physical Interface	Logical Interface	Status Description
ge-3/2/0 Admin Up Link Down	ge-3/2/0.0 Admin Up Link Down	This interface has both the physical link and the logical interface down.

SEE ALSO

Monitoring Fast Ethernet and Gigabit Ethernet

Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface

IN THIS SECTION

- Purpose | 507
- Action | **507**
- Meaning | **509**

Purpose

To display the status of a specific Fast Ethernet or Gigabit Ethernet interface when you need to investigate its status further, use the following Junos OS CLI operational mode command:

Action

user@host> show interfaces (fe-fpc/pic/port | ge-fpc/pic/port)

Sample Output 1

The following sample output is for a Fast Ethernet interface with the physical link up:

```
user@host> show interfaces fe-2/1/0
Physical interface: fe-2/1/0, Enabled, Physical link is Up
 Interface index: 31, SNMP ifIndex: 35
 Description: customer connection
 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: 00:90:69:86:71:1b, Hardware address: 00:90:69:86:71:1b
 Input rate : 25768 bps (11 pps), Output rate: 1576 bps (3 pps)
 Active alarms : None
 Active defects : None
 Logical interface fe-2/1/0.0 (Index 2) (SNMP ifIndex 43)
   Flags: SNMP-Traps, Encapsulation: ENET2
   Protocol inet, MTU: 1500, Flags: Is-Primary
     Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.116.151.218/29, Local: 10.119.115.217
       Broadcast: 10.116.151.225
```

Sample Output 2

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. This means that the physical link is healthy and can pass packets. Further down the sample output, look for active alarms and defects. If you see active alarms or defects, to further diagnose the problem, see Step 3, "Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface" on page 503, to display more extensive information about the Fast Ethernet interface and the physical interface that is down.

Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface

IN THIS SECTION

Purpose | **509**

Action | **509**

Meaning | 511

Purpose

To display extensive status information about a specific Fast Ethernet or 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 Fast 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
                                                     Transmit
MAC statistics:
                                     Receive
                                   439703575
                                                    177452093
  Total octets
  Total packets
                                     1866532
                                                      1642916
  Unicast packets
                                      972137
                                                      1602563
  Broadcast packets
                                          30
                                                         2980
  Multicast packets
                                      894365
                                                        37373
                                           0
                                                            0
  CRC/Align errors
  FIFO errors
                                           0
                                                            0
  MAC control frames
                                                            0
                                                            0
  MAC pause frames
                                           0
  Oversized frames
                                           0
  Jabber frames
  Fragment frames
                                           0
                                           0
  VLAN tagged frames
  Code violations
                                           0
Filter statistics:
                                     1866532
  Input packet count
  Input packet rejects
                                           0
   Input DA rejects
                                       503674
   Input SA rejects
                                            0
  Output packet count
                                                      1642916
  Output packet pad count
```

```
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. See Table 4 for a description of errors to look for.

Table 158: Errors to Look For

Error	Meaning
Policed discards	Discarded frames that were not recognized or were not of interest.
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 Fast Ethernet or 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.

Table 158: Errors to Look For (Continued)

Error	Meaning
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 Fast Ethernet or Gigabit Ethernet interface and diagnose the Fast Ethernet or Gigabit Ethernet media accordingly. See "Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters" on page 566 for an explanation of Fast Ethernet and Gigabit Ethernet alarms.

Table 5 lists and describes some MAC statistics errors to look for.

Table 159: 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.	
Jabber frames	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.	
	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.	
Fragment frames	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 an alignment error. 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.	

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 160: Autonegotiation Information

Autonegotiation Field Information	Explanation
Negotiation status: Incomplete	The Negotiation status field shows Incomplete when the Ethernet interface has the speed or link mode configured.
Negotiation status: No autonegotiation	The Negotiation status field shows No autonegotiation when the remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
Negotiation status: Complete Link partner status: OK	The Negotiation status field shows Complete and the Link partner field shows OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process completes successfully.
Link partner: Half-duplex	The Link partner field can be Full-duplex or Half-duplex depending on the capability of the attached Ethernet device.
Flow control: Symmetric/ asymmetric	The Flow control field displays the types of flow control supported by the remote Ethernet device.

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface

IN THIS SECTION

• Purpose | **514**

Action | **514**

Meaning | 515

Purpose

To monitor statistics for a Fast Ethernet or 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)



CAUTION: 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 Fast Ethernet interface:

usardhast> manitar interf	ace fe-2/1/0			
user@host> monitor interface fe-2/1/0 Interface: fe-2/1/0, Enabled, Link is Up				
·				
Encapsulation: Ethernet, S	speed: Taallibbs		C D.14.	
Traffic statistics:	000550004040	(4.4000)	Current Delta	
Input bytes:	282556864218	·	[40815]	
Output bytes:	42320313078	(384 bps)	[890]	
Input packets:	739373897	(11 pps)	[145]	
Output packets:	124798688	(1 pps)	[14]	
Error statistics:				
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 statistic	cs:		
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.

NOTE: 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 statistics in the second column are the cumulative statistics since the last time they were cleared using the clear interfaces statistics *interface-name* command. The statistics in the third column are the cumulative statistics since the monitor interface *interface-name* command was executed.

If the input errors are increasing, verify the following:

- 1. Check the cabling to the router and have the carrier verify the integrity of the line. To verify the integrity of the cabling, make sure that you have the correct cables for the interface port. Make sure you have single-mode fiber cable for a single-mode interface and multimode fiber cable for a multimode interface.
- 2. For a fiber-optic connection, measure the received light level at the receiver end and make sure that it is within the receiver specification of the Ethernet interface. See "Fiber-Optic Ethernet Interface Specifications" on page 515 for the fiber-optic Ethernet interface specifications.
- **3.** Measure the transmit light level on the Tx port to verify that it is within specification. See "Fiber-Optic Ethernet Interface Specifications" on page 515 for the optical specifications.

Fiber-Optic Ethernet Interface Specifications

Table 161 on page 515 shows the specifications for fiber-optic interfaces for Juniper Networks routers.

Table 161: Fiber-Optic Ethernet Interface Specifications

Fiber-Optic Ethernet Interface	Length	Wavelength	Average Launch Power	Receiver Saturation	Receiver Sensitivity
Gigabit Ethernet					
Duplex SC connector					

Table 161: Fiber-Optic Ethernet Interface Specifications (Continued)

Fiber-Optic Ethernet Interface	Length	Wavelength	Average Launch Power	Receiver Saturation	Receiver Sensitivity
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 1012) 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
Fast Ethernet 8-Port					
FX optical interface with MT- RJ connector	1.24-mile 2-km reach on 62.5/125- micrometer MMF	1270 to 1380 nm	-20 to -14 dBm	-14 dBm	-34 dBm

SEE ALSO

Ethernet Interfaces User Guide for Routing Devices

Fiber-Optic Cable Signal Loss, Attenuation, and Dispersion

Calculating Power Budget and Power Margin for Fiber-Optic Cables

Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

IN THIS SECTION

- Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 517
- Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface | 519
- Create a Loopback | 520
- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 523
- Configure a Static Address Resolution Protocol Table Entry | 527
- Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 532
- Ping the Fast Ethernet or Gigabit Ethernet Interface | 533
- Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 535
- Diagnose a Suspected Circuit Problem | 538

Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

IN THIS SECTION

- Purpose | **517**
- Action | 517

Purpose

To use loopback testing to isolate Fast Ethernet and Gigabit Ethernet interface problems.

Action

Table 162 on page 518 provides links and commands for using loopback testing for Fast Ethernet and Gigabit Ethernet interfaces.

Table 162: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Tasks	Command or Action			
"Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface" on page 519				
1. "Create a Loopback" on page 520				
1. a. "Create a Physical Loopback for a Fiber-Optic Interface" on page 520	Connect the transmit port to the receive port.			
1. a. "Create a Loopback Plug for an RJ-45 Ethernet Interface" on page 520	Cross pin 1 (TX+) and pin 3 (RX+) together, and pin 2 (TX-) and pin 6 (RX-) together.			
1. a. "Configure a Local Loopback" on page 517	[edit interfaces <i>interface-name</i> (fastether-options gigether-options)] set loopback show commit			
1. "Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up" on page 523	show interfaces (fe-fpc/pic/port ge-fpc/pic/port)			
"Configure a Static Address Resolution Protocol Table Entry" on page 527	show interfaces ge-fpc/pic/port [edit interfaces interface-name unit logical-unit-number family inet address address] set arp ip-address mac mac-address show commit run show arp no-resolve			
"Clear Fast Ethernet or Gigabit Ethernet Interface Statistics" on page 532	clear interfaces statistics fe-fpc/pic/port / ge-fpc/pic/port			
1. "Ping the Fast Ethernet or Gigabit Ethernet Interface" on page 533	ping remote-IP-address bypass-routing interface (fe-fpc/pic/port / ge-fpc/pic/port count 100 rapid			
"Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics" on page 535	<pre>show interfaces (fe-fpc/pic/port ge-fpc/pic/ port) extensive</pre>			

Table 162: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces (Continued)

Tasks	Command or Action
"Diagnose a Suspected Circuit Problem" on page 538	Perform Steps 2 through 8 from "Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface" on page 519.

Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface

IN THIS SECTION

• Problem | **519**

Solution | 519

Problem

Description

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" on page 520
- "Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up" on page 523
- "Configure a Static Address Resolution Protocol Table Entry" on page 527
- "Clear Fast Ethernet or Gigabit Ethernet Interface Statistics" on page 532
- "Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics" on page 535

Create a Loopback

IN THIS SECTION

- Create a Physical Loopback for a Fiber-Optic Interface | 520
 - Create a Loopback Plug for an RJ-45 Ethernet Interface | 520
- Configure a Local Loopback | 521

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

Action

To create a physical loopback at the port, connect the transmit port to the receive port using a known good fiber cable.

NOTE: Make sure you use single-mode fiber for a single-mode port and multimode fiber for a multimode port.

Meaning

When you create and then test a physical loopback, you are testing the transmit and receive ports 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

Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 517

Create a Loopback Plug for an RJ-45 Ethernet Interface

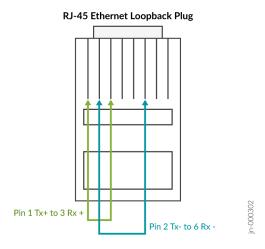
Action

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

Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 517

Configure a Local Loopback

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 (fastether-options | gigether-options)
```

2. Configure the local loopback:

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit interfaces fe-1/0/0 fastether-options]
user@host# show
loopback;
```

4. Commit the change:

```
user@host# commit
```

For example:

```
[edit interfaces fe-1/0/0 fastether-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* (fastether-options | gigether-options] hierarchy level, places the interface into local loopback mode.

NOTE: Remember to delete the loopback statement after completing the test.

SEE ALSO

Create a Loopback | 520

RELATED DOCUMENTATION

Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 523

Configure a Static Address Resolution Protocol Table Entry | 527

Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 532

Ping the Fast Ethernet or Gigabit Ethernet Interface | 533

Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 535

Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up

IN THIS SECTION

- Purpose | **523**
- Action | **524**
- Meaning | 527

Purpose

Display the status of the Fast Ethernet or 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 Fast Ethernet or 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
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:
                                 0
  Input bytes :
                                                    0 bps
  Output bytes :
                                 0
                                                    0 bps
  Input packets:
                                                    0 pps
  Output packets:
                                 0
                                                    0 pps
  IPv6 transit statistics:
   Input bytes :
                                 0
   Output bytes :
                                 0
   Input packets:
                                 0
   Output packets:
 Dropped traffic statistics due to STP State:
  Input bytes :
  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
Oueue counters:
                      Queued packets Transmitted packets
                                                                Dropped packets
  0 best-effort
                                    0
  1 expedited-fo
                                    0
                                                         0
                                                                               0
  2 assured-forw
                                    0
                                                                               0
  3 network-cont
                                    0
                                                         0
                                                                               0
Queue number:
                      Mapped forwarding classes
                      best-effort
  1
                      expedited-forwarding
  2
                      assured-forwarding
  3
                      network-control
Active alarms : None
Active defects : None
MAC statistics:
                                                      Transmit
                                      Receive
  Total octets
                                                             0
                                            0
  Total packets
                                            0
                                                             0
  Unicast packets
                                            0
                                                             0
  Broadcast packets
                                                             0
  Multicast packets
                                                             0
                                            0
  CRC/Align errors
                                            0
                                                             0
  FIFO errors
                                                             0
  MAC control frames
                                                             0
                                            0
  MAC pause frames
                                            0
                                                             0
  Oversized frames
  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
  Output packet count
                                                             0
                                                             0
  Output packet pad count
  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
      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
                                                                0
                                                95
                                                                       low
                                                                              none
  3 network-control
                            5
                                    50000000
                                                 5
                                                                       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
                : 0 bps (0 pps), Output rate: 0 bps (0 pps)
 Input rate
 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 163 on page 527 presents problem situations and actions for a physical link that is down.

Table 163: 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

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 Fast 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
                : None
 Link flags
 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:
   Output packets:
 Dropped traffic statistics due to STP State:
  Input bytes :
  Output bytes :
                                     0
  Input packets:
                                     0
  Output packets:
```

```
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
 1 expedited-fo
                                   0
                                                         0
                                    0
                                                         0
  2 assured-forw
  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
                                         1280
 Unicast packets
                                                          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
  Oversized frames
                                            0
  Jabber frames
                                            0
  Fragment frames
                                            0
  VLAN tagged frames
                                            0
  Code violations
                                            0
Filter statistics:
  Input packet count
                                         1281
                                            0
  Input packet rejects
 Input DA rejects
  Input SA rejects
                                            0
                                                          1281
  Output packet count
  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
                                                                0
                           5
                                    50000000
                                                 5
                                                                       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
                                   1281
  Input packets:
  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
                                      0
  Output packets:
                                                           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 :
```

```
Bytes permitted by policy :
                                     107520
  Connections established :
                                     20
Flow Output statistics:
  Multicast packets :
  Bytes permitted by policy :
                                     107520
Flow error statistics (Packets dropped due to):
  Address spoofing:
  Authentication failed:
  Incoming NAT errors:
  Invalid zone received packet:
                                     0
  Multiple user authentications:
  Multiple incoming NAT:
  No parent for a gate:
  No one interested in self packets: 0
  No minor session:
  No more sessions:
                                     0
  No NAT gate:
  No route present:
                                     11
  No SA for incoming SPI:
  No tunnel found:
  No session for a gate:
  No zone or NULL zone binding
                                     0
  Policy denied:
  Security association not active:
 TCP sequence number out of window: 0
  Syn-attack protection:
  User authentication errors:
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:

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	9 10.108.120.2	ge-4/0/6.0		permanent
52:54:00:7d:33:4d	10.204.128.35	fxp0.0		none
52:54:00:65:11:50	0 10.204.128.36	fxp0.0		none
52:54:00:da:30:82	2 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	7 10.204.128.46	fxp0.0		none

Meaning

The sample output is for Step 1 through Step 6 and shows that a static ARP entry was configured on Gigabit Ethernet interface ge-4/0/6.

Clear Fast Ethernet or Gigabit Ethernet Interface Statistics

Purpose

You can reset the Fast Ethernet and 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:

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 Fast Ethernet or Gigabit Ethernet Interface

IN THIS SECTION

- Purpose | **533**
- Action | **533**
- Meaning | 534

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:

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
                                                    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
                                                    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
                                                     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
                                                    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
                                                    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 the frames from the ping test. The MAC address used is the same as the physical address of the port being tested because this allows the port to accept the frames from the ping test. As the packet is looped over the link, you expect to receive a TLL exceeded message for each ping sent. These messages are generated because the ping packets are repeatedly looped between the router and the physical loopback. When the packet is sent to the other end of the link, which does not exist, the loopback returns the packet back to the same interface, where it is again subjected to the Packet Forwarding Engine fabric for routing. After the route lookup, the TTL is decremented, and the packet is again sent out of the looped interface. This process repeats until the packed is either lost, or the TLL expires with subsequent TTL expired message displayed. Should any errors occur, the packet is discarded and a time-out error is displayed, rather than the expected TTL expired message. Note that the default TTL for ICMP echo packets in Junos OS is 64. This means a given test packet must be successfully sent and received 63 times before a TTL expired message can be generated. You can alter the TTL value to adjust the tolerance for loss, for example, a value of 255 is the most demanding test because now the packet must be sent and received error free 254 times.

Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics

IN THIS SECTION

- Purpose | **535**
- Action | **535**
- Meaning | 537

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:

Sample Output

```
user@router> show interfaces ge-4/0/6 extensive
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 pps
 IPv6 transit statistics:
 Input bytes :
                                    0
 Output bytes :
 Input packets:
                                    0
  Output packets:
Dropped traffic statistics due to STP State:
 Input bytes :
 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
                      Queued packets Transmitted packets
Oueue counters:
                                                               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
Oueue number:
                      Mapped forwarding classes
  0
                      best-effort
                      expedited-forwarding
  1
  2
                      assured-forwarding
                      network-control
Active alarms : None
Active defects : None
MAC statistics:
                                     Receive
                                                     Transmit
  Total octets
                                           0
                                                             0
                                                             0
 Total packets
                                           0
  Unicast packets
                                           0
                                                             0
  Broadcast packets
                                           0
                                                             0
                                           0
                                                             0
  Multicast packets
                                           0
                                                             0
  CRC/Align errors
  FIFO errors
```

```
MAC control frames
                                            0
                                                             0
  MAC pause frames
                                                             0
                                            0
  Oversized frames
                                            0
  Jabber frames
                                            0
  Fragment frames
  VLAN tagged frames
                                            0
  Code violations
                                            0
Filter statistics:
  Input packet count
                                            0
  Input packet rejects
                                            0
  Input DA rejects
  Input SA rejects
                                            0
  Output packet count
                                                             0
  Output packet pad count
                                                             0
  Output packet error count
  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
                                                            Buffer Priority
  CoS transmit queue
                                   Bandwidth
                                                                               Limit
                            %
                                          bps
                                                  %
                                                              usec
  0 best-effort
                           95
                                   950000000
                                                 95
                                                                 0
                                                                        low
                                                                                none
  3 network-control
                                     50000000
                                                  5
                                                                        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 | **538**
- Action | 538

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 Fast Ethernet or Gigabit Ethernet Interface" on page 519. Keep in mind that 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.

Performing Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces

IN THIS SECTION

- Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 539
- Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | 540
- Create a Loopback | 541
- Verify That the Gigabit Ethernet Interface Is Up | 544

- Configure a Static Address Resolution Protocol Table Entry | 549
- Clear Gigabit Ethernet Interface Statistics | 551
- Ping the Gigabit Ethernet Interface | 551
- Check for Gigabit Ethernet Interface Error Statistics | 553
- Diagnose a Suspected Circuit Problem | 557

Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces

IN THIS SECTION

- Purpose | **539**
- Action | 539

Purpose

To use loopback testing to isolate 10, 40, and 100 Gigabit Ethernet interface problems.

Action

Table 164 on page 539 provides links and commands for using loopback testing for 10, 40, and 100 Gigabit Ethernet interfaces.

Table 164: Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces

Tasks	Command or Action
"Diagnose a Suspected Hardware Problem with a Gigabit Ethernet I	nterface" on page 540
1. "Create a Loopback" on page 541	
"Create a Physical Loopback for a Fiber-Optic Interface" on page 520	Connect the transmit port to the receive port.

Table 164: Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces *(Continued)*

Tasks	Command or Action
1. "Configure a Local Loopback" on page 541	<pre>[edit interfaces interface-name (gigether- options)] set loopback show commit</pre>
"Verify That the Gigabit Ethernet Interface Is Up" on page 544	<pre>show interfaces (xe-fpc/pic/port et-fpc/pic/ port)</pre>
"Configure a Static Address Resolution Protocol Table Entry" on page 549	show interfaces (xe-fpc/pic/port et-fpc/pic/port)[edit interfaces interface-name unit logical-unit-number family inet address address] set arp ip-address mac mac-address show commit run show arp no-resolve
1. "Clear Gigabit Ethernet Interface Statistics" on page 551	<pre>clear interfaces statistics xe-fpc/pic/port / et-fpc/pic/port</pre>
1. "Ping the Gigabit Ethernet Interface" on page 551	<pre>ping remote-IP-address bypass-routing interface (xe-fpc/pic/port / et-fpc/pic/port count 100</pre>
"Check for Gigabit Ethernet Interface Error Statistics" on page 553	show interfaces (xe- $fpc/pic/port$ et- $fpc/pic/port$) extensive
"Diagnose a Suspected Circuit Problem" on page 557	Perform Steps 2 through 8 from "Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface" on page 519 .

Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface

IN THIS SECTION

Problem | **541**

Solution | 541

Problem

Description

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" on page 541
- "Verify That the Gigabit Ethernet Interface Is Up" on page 544
- "Configure a Static Address Resolution Protocol Table Entry" on page 549
- "Clear Gigabit Ethernet Interface Statistics" on page 551
- "Check for Gigabit Ethernet Interface Error Statistics" on page 553

Create a Loopback

IN THIS SECTION

- Create a Physical Loopback for a Fiber-Optic Interface | 541
- Configure a Local Loopback | 542

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

Action

To create a physical loopback at the port, connect the transmit port to the receive port using a known good fiber cable.

NOTE: Make sure you use single-mode fiber for a single-mode port and multimode fiber for a multimode port.

Meaning

When you create and then test a physical loopback, you are testing the transmit and receive ports 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

Verify That the Gigabit Ethernet Interface Is Up | 544

Configure a Static Address Resolution Protocol Table Entry | 549

Configure a Local Loopback

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@router# edit interfaces interface-name gigether-options

2. Configure the local loopback:

[edit interfaces interface-name gigether-options

user@router# set

loopback

3. Verify the configuration:

user@router# show

For example:

```
[edit interfaces xe-2/0/0 gigether-options]
user@router# show
loopback;
```

4. Commit the change:

```
user@router# commit
```

For example:

```
[edit interfaces xe-2/0/0 gigether-options]
user@router# 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.

NOTE: Remember to delete the loopback statement after completing the test.

SEE ALSO

Verify That the Gigabit Ethernet Interface Is Up | 544

Configure a Static Address Resolution Protocol Table Entry | 549

RELATED DOCUMENTATION

Verify That the Gigabit Ethernet Interface Is Up | 544

Configure a Static Address Resolution Protocol Table Entry | 549

Clear Gigabit Ethernet Interface Statistics | 551

Ping the Gigabit Ethernet Interface | 551

Check for Gigabit Ethernet Interface Error Statistics | 553

Verify That the Gigabit Ethernet Interface Is Up

IN THIS SECTION

- Purpose | **544**
- Action | **544**
- Meaning | 548

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@router> show interfaces (xe-fpc/pic/port |
et-fpc/pic/port)
```

Sample Output

```
user@router# show interfaces xe-2/0/0 extensive
Physical interface: xe-2/0/0, Enabled, Physical link is Up
Interface index: 187, SNMP ifIndex: 591, Generation: 190
Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loop Detect PDU Error: None,
MAC-REWRITE Error: None, Loopback: Local, Source filtering: Disabled, Flow control: Enabled,
Speed Configuration: Auto
Pad to minimum frame size: Disabled
Device flags: Present Running Loop-Detected
Interface flags: SNMP-Traps Internal: 0x4000
Link flags: None
CoS queues: 8 supported, 8 maximum usable queues
```

```
Schedulers
                : 0
                : Up 4000 ms, Down 0 ms
 Hold-times
 Damping
                 : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
 Current address: d8:18:d3:b3:6d:ea, Hardware address: d8:18:d3:b3:6d:ea
 Last flapped : 2019-04-20 17:13:55 PDT (13w4d 21:42 ago)
 Statistics last cleared: 2019-07-25 14:55:21 PDT (00:01:01 ago)
  Traffic statistics:
  Input bytes :
                                 537600
                                                           0 bps
  Output bytes :
                                 539600
                                                           0 bps
  Input packets:
                                   6400
                                                           0 pps
  Output packets:
                                   6400
                                                           0 pps
  IPv6 transit statistics:
  Input bytes :
  Output bytes :
                                      0
  Input packets:
                                      0
  Output packets:
  Dropped traffic statistics due to STP State:
  Input bytes :
  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
                                  6400
                                                       6400
    1
                                     0
                                                          0
                                                                               0
   2
                                     0
                                                          0
                                                                               0
   3
  Queue number:
                        Mapped forwarding classes
                        best-effort
    0
                        expedited-forwarding
   1
    2
                        assured-forwarding
    3
                        network-control
 Active alarms : None
  Active defects : None
```

PCS statistics		Secon	ds					
Bit errors			0					
Errored blocks			0					
MAC statistics:		Rece	ive	Т	ransmit			
Total octets		652	800		652800			
Total packets			400		6400			
Unicast packets			400		6400			
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					
Total errors			0		0			
Filter statistics:								
Input packet count		6	400					
Input packet rejects			0					
Input DA rejects			0					
Input SA rejects			0					
Output packet count			-		6400			
Output packet pad co	unt				0			
Output packet error					0			
CAM destination filt		M source	filte	-s· 0	ŭ			
Packet Forwarding Engi				J. U				
Destination slot: 0	_							
CoS information:	(0,00)							
Direction : Output								
CoS transmit queue		Bandwi	dth		Buffe	er Priori	t.v	Limit
4	%		bps	%	use		-3	
0 best-effort	95	9500000		95	450		OW	none
3 network-control	5	500000		5			OW	none
Preclassifier statisti		20000		-		_		
Traffic Class	Received	Packets	Trans	smitted	Packets	Dron	ped	Packets
real-time		0			0	5.00	₁ u	0
Teal - Clille					6400			0
		6400						
network-control		6400 0						
		6400 0			0			0

```
Interface transmit statistics: Disabled
 Logical interface xe-2/0/0.0 (Index 353) (SNMP ifIndex 599) (Generation 175)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
   Traffic statistics:
    Input bytes :
                                   537600
    Output bytes :
                                   539000
    Input packets:
                                     6400
    Output packets:
                                     6400
    Local statistics:
    Input bytes :
                                        0
    Output bytes :
                                     9800
    Input packets:
                                        0
    Output packets:
                                      100
    Transit statistics:
    Input bytes :
                                                             0 bps
                                   537600
    Output bytes :
                                   529200
                                                             0 bps
    Input packets:
                                     6400
                                                             0 pps
    Output packets:
                                     6300
                                                             0 pps
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
   Generation: 206, 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:
146
    Protocol multiservice, MTU: Unlimited, Generation: 207, Route table: 0
      Policer: Input: __default_arp_policer__
```

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 et-*fpcl picl port* command when the physical link is down:

```
user@router> show interfaces et-3/0/1
Physical interface: et-3/0/1, Enabled, Physical link is Down
Interface index: 620, SNMP ifIndex: 564
```

```
Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 40Gbps, BPDU Error: None, Loop Detect
PDU Error: None, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Pad to minimum frame size: Disabled
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Schedulers
 Current address: 54:e0:32:71:e3:e2, Hardware address: 54:e0:32:71:e3:e2
 Last flapped : 2019-07-05 09:10:02 PDT (3d 14:46 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
                                         2
   Errored blocks
                                         6
 Ethernet FEC Mode :
                                       NONE
 Ethernet FEC statistics
                                     Errors
   FEC Corrected Errors
                                          0
   FEC Uncorrected Errors
                                          0
   FEC Corrected Errors Rate
                                          0
   FEC Uncorrected Errors Rate
 Interface transmit statistics: Disabled
```

Meaning

The sample output shows that the physical link is down and there are active alarms and defects.

Table 165 on page 548 presents problem situations and actions for a physical link that is down.

Table 165: 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.

Table 165: Problems and Solutions for a Physical Link That Is Down (Continued)

Problem	Action
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

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:

1. Find the Media Access Control (MAC) address for the Gigabit Ethernet interface:

user@router# run show interfaces xe-2/0/0 extensive | match "Current address"

Current address: d8:18:d3:b3:6d:ea, Hardware address: d8:18:d3:b3:6d:ea

2. In configuration mode, go to the following hierarchy level:

3. Configure the static ARP entry:

```
user@router# set arp ip-address mac mac-address
```

4. Commit the configuration:

```
user@router# commit
```

5. Verify that the static ARP entry is installed:

```
[edit interfaces xe-2/0/0.0 unit 0 family inet address 10.108.120.1/30] user@router#run show arp no-resolve
```

C Address Address	Interface	Flags
:01:00:00:00:05 10.0.0.5	em1.0	none
:00:5e:00:01:01 10.85.175.1	fxp0.0	none
:b1:22:0a:6e:00 10.85.175.2	fxp0.0	none
:07:ca:57:d7:a0 10.85.175.3	fxp0.0	none
:a0:a5:c2:06:e2 10.85.175.4	fxp0.0	none
:18:d3:b3:6d:ea 10.108.120.2	xe-2/0/0.0	permanent
:01:00:00:00:05 128.0.0.5	em1.0	none
:01:00:00:00:05 128.0.0.6	em1.0	none
:00:00:00:00:10 128.0.0.16	em0.0	none
:00:00:00:00:12 128.0.0.18	em0.0	none
:00:00:00:00:17 128.0.0.23	em0.0	none
:00:00:00:00:1a 128.0.0.26	em0.0	none
tal entries: 12		

Meaning

The sample output is for Step 1 through Step 6 and shows that a static ARP entry was configured on Gigabit Ethernet interface xe-2/0/0.0.

Clear Gigabit Ethernet Interface Statistics

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:

Sample Output

```
user@router> clear interfaces statistics xe-2/0/0
user@router>
```

Meaning

This command clears the interface statistics counters for the Gigabit Ethernet interface only.

Ping the Gigabit Ethernet Interface

IN THIS SECTION Purpose | 551 Action | 552 Meaning | 552

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:

Sample Output

```
user@router> ping 10.108.120.2 bypass-routing interface xe-2/0/0 count 100 rapid
ping 10.108.120.2 bypass-routing interface xe-2/0/0 count 100 rapid
(snip)
Vr HL TOS Len ID Flg off TTL Pro cks
                                            Src
                                                     Dst
4 5 00 0054 6a14 0 0000 36 01 15ba 10.108.120.1 10.108.120.2
36 bytes from 10.108.120.1: Redirect Host(New addr: 10.108.120.2)
Vr HL TOS Len ID Flg off TTL Pro cks
                                                     Dst
4 5 00 0054 6a14 0 0000 35 01 16ba 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
                                                     Dst
4 5 00 0054 6a14 0 0000 01 01 4aba 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
                                                     Dst
4 5 00 0054 6a76 0 0000 01 01 4a58 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
                                                     Dst
4 5 00 0054 6aa7 0 0000 01 01 4a27 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
                                                     Dst
4 5 00 0054 6ae3 0 0000 01 01 49eb 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
                                                     Dst
4 5 00 0054 6b0a
                   0 0000 01 01 49c4 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 the frames from the ping test. The MAC address used is the same as the physical address of the port being tested because this allows the port to accept the frames from the ping test. As the packet is looped over the link, you expect to receive a TLL exceeded message for each ping sent. These messages are

generated because the ping packets are repeatedly looped between the router and the physical loopback. When the packet is sent to the other end of the link, which does not exist, the loopback returns the packet back to the same interface, where it is again subjected to the Packet Forwarding Engine fabric for routing. After the route lookup, the TTL is decremented, and the packet is again sent out of the looped interface. This process repeats until the packed is either lost, or the TLL expires with subsequent TTL expired message displayed. Should any errors occur, the packet is discarded and a time-out error is displayed, rather than the expected TTL expired message. Note that the default TTL for ICMP echo packets in Junos OS is 64. This means a given test packet must be successfully sent and received 63 times before a TTL expired message can be generated. You can alter the TTL value to adjust the tolerance for loss, for example, a value of 255 is the most demanding test because now the packet must be sent and received error free 254 times.

Check for Gigabit Ethernet Interface Error Statistics

IN THIS SECTION

- Purpose | 553
- Action | **553**
- Meaning | 557

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:

Sample Output

```
user@router> show interfaces xe-2/0/0 extensive

Physical interface: xe-2/0/0, Enabled, Physical link is Up

Interface index: 187, SNMP ifIndex: 591, Generation: 190

Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
```

```
None, Loop Detect PDU Error: None,
  MAC-REWRITE Error: None, Loopback: Local, Source filtering: Disabled, Flow control: Enabled,
Speed Configuration: Auto
 Pad to minimum frame size: Disabled
 Device flags : Present Running Loop-Detected
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags
                : None
 CoS queues
                : 8 supported, 8 maximum usable queues
 Schedulers
                : Up 4000 ms, Down 0 ms
 Hold-times
 Damping
                : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
 Current address: d8:18:d3:b3:6d:ea, Hardware address: d8:18:d3:b3:6d:ea
 Last flapped : 2019-04-20 17:13:55 PDT (13w4d 21:40 ago)
 Statistics last cleared: 2019-07-25 14:49:21 PDT (00:04:41 ago)
 Traffic statistics:
  Input bytes :
                                537600
                                                           0 bps
  Output bytes :
                                539600
                                                          0 bps
  Input packets:
                                  6400
                                                          0 pps
  Output packets:
                                  6400
                                                          0 pps
  IPv6 transit statistics:
  Input bytes :
                                     0
  Output bytes :
                                     0
  Input packets:
  Output packets:
  Dropped traffic statistics due to STP State:
  Input bytes :
  Output bytes :
                                     0
  Input packets:
                                     0
  Output packets:
 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
                                  6400
                                                      6400
                                                                               0
                                                                               0
   1
                                    0
                                                          0
    2
```

3	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				
PCS statistics	Seconds			
Bit errors	0			
Errored blocks	0			
MAC statistics:	Receive	Transmit		
Total octets	652800	652800		
Total packets	6400	6400		
Unicast packets	6400	6400		
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			
Total errors	0	0		
Filter statistics:				
Input packet count	6400			
Input packet reject	s 0			
Input DA rejects	0			
Input SA rejects	0			
Output packet count		6400		
Output packet pad c	ount	0		
Output packet error	count	0		
CAM destination fil	ters: 0, CAM source filters	s: 0		
Packet Forwarding Eng	ine configuration:			
Destination slot: 0	(0x00)			
CoS information:				
Direction : Output				
CoS transmit queue	Bandwidth	Buf	fer Priority	Limit
	% bps	% u	sec	
0 best-effort	95 9500000000	95	0 low	none

```
3 network-control
                                    500000000
                                                                         low
                             5
                                                  5
                                                                                none
 Preclassifier statistics:
   Traffic Class
                        Received Packets Transmitted Packets
                                                                      Dropped Packets
    real-time
                                        0
                                                               0
   network-control
                                     6400
                                                            6400
                                                                                    0
   best-effort
                                                                                    0
                                        0
                                                               0
 Link Degrade :
    Link Monitoring
                                      : Disable
 Interface transmit statistics: Disabled
 Logical interface xe-2/0/0.0 (Index 353) (SNMP ifIndex 599) (Generation 175)
   Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
   Traffic statistics:
    Input bytes :
                                  537600
    Output bytes :
                                   539000
    Input packets:
                                    6400
    Output packets:
                                    6400
    Local statistics:
    Input bytes :
                                        0
    Output bytes :
                                    9800
    Input packets:
                                        0
    Output packets:
                                     100
   Transit statistics:
    Input bytes :
                                  537600
                                                             0 bps
    Output bytes :
                                   529200
                                                             0 bps
    Input packets:
                                    6400
                                                             0 pps
    Output packets:
                                    6300
                                                             0 pps
    Protocol inet, MTU: 1500
   Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop
cnt: 0
    Generation: 206, 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:
146
   Protocol multiservice, MTU: Unlimited, Generation: 207, Route table: 0
      Policer: Input: __default_arp_policer__
```

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 | **557**
- Action | 557

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" on page 540. Keep in mind that 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.

NOTE: This document is applicable for 1Gb, 10Gb, 40Gb, and 100Gb interfaces.

Configuring Interface Diagnostics Tools to Test the Physical Layer Connections

IN THIS SECTION

- Configuring Loopback Testing | 558
- Configuring BERT Testing | 560
- Starting and Stopping a BERT Test | 565

Configuring Loopback Testing

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, Fast Ethernet, Gigabit Ethernet, E1, E3, NxDSO, serial, SONET/SDH, T1, and T3.

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.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own Data Link Layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own Data Link Layer packets, you can assume that the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

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. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. 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. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 166 on page 559 shows the loopback modes supported on the various interface types.

Table 166: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	Configuring Ethernet Loopback Capability
Circuit Emulation E1	Local and remote	Configuring E1 Loopback Capability
Circuit Emulation T1	Local and remote	Configuring T1 Loopback Capability
E1 and E3	Local and remote	Configuring E1 Loopback Capability and Configuring E3 Loopback Capability
Λ ∕xDS0	Payload	Configuring NxDS0 IQ and IQE Interfaces, Configuring T1 and NxDS0 Interfaces, Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode), Configuring Fractional E1 IQ and IQE Interfaces, and Configuring Channelized T3 IQ Interfaces
Serial (V.35 and X.21)	Local and remote	Configuring Serial Loopback Capability

Table 166: 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
SONET/SDH	Local and remote	Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External
T1 and T3	Local, payload, and remote	Configuring T1 Loopback Capability and Configuring T3 Loopback Capability See also Configuring the T1 Remote Loopback Response

To configure loopback testing, include the loopback statement:

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* ds0-options]
- [edit interfaces *interface-name* e1-options]
- [edit interfaces *interface-name* e3-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]
- [edit interfaces *interface-name* serial-options]
- [edit interfaces *interface-name* sonet-options]
- [edit interfaces *interface-name* t1-options]
- [edit interfaces *interface-name* t3-options]

Configuring BERT Testing

To configure BERT:

• 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.

• 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).

• 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.

NOTE: The four-port E1 PIC supports only the following algorithms:

```
pseudo-2e11-o152 Pattern is 2^11 -1 (per 0.152 standard)
pseudo-2e15-o151 Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151 Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e23-o151 Pattern is 2^23 (per 0.151 standard)
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```
all-ones-repeating
                     Repeating one bits
all-zeros-repeating Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e11-o152
                     Pattern is 2^11 -1 (per 0.152 standard)
                     Pattern is 2<sup>15</sup> - 1 (per 0.151 standard)
pseudo-2e15-o151
pseudo-2e20-o151
                     Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7
                     Pattern is 2<sup>7</sup> - 1
pseudo-2e9-o153
                     Pattern is 2<sup>9</sup> - 1 (per 0.153 standard)
                     1 bit in 4 is set
repeating-1-in-4
                     1 bit in 8 is set
repeating-1-in-8
                      3 bits in 24 are set
repeating-3-in-24
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: The IQE PICs support only the following algorithms:

```
all-ones-repeating Repeating one bits
all-zeros-repeating Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e9-o153 Pattern is 2^9 -1 (per 0.153 (511 type) standard)
pseudo-2e11-o152 Pattern is 2^11 -1 (per 0.152 and 0.153 (2047 type) standards)
pseudo-2e15-o151 Pattern is 2^15 -1 (per 0.151 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)
pseudo-2e23-o151 Pattern is 2^23 -1 (per 0.151 standard)
repeating-1-in-4 1 bit in 4 is set
repeating-1-in-8 1 bit in 8 is set
repeating-3-in-24 3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

```
all-ones-repeating
                               Repeating one bits
all-zeros-repeating
                               Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros
                               Alternating ones and zeros
                              1 bit in 4 is set
repeating-1-in-4
repeating-1-in-8
                               1 bit in 8 is set
repeating-3-in-24
                               3 bits in 24 are set
                               Pattern is 2<sup>9</sup> - 1 (per 0.153 standard)
pseudo-2e9-o153
pseudo-2e11-o152
                               Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151
                               Pattern is 2<sup>15</sup> - 1 (per 0.151 standard)
                               Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o151
                               Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e20-o153
                               Pattern is 2^23 (per 0.151 standard)
pseudo-2e23-o151
```

Table 167 on page 563 shows the BERT capabilities for various interface types.

Table 167: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0-11)	_	Limited algorithms

Table 167: BERT Capabilities by Interface Type (Continued)

Interface	T1 BERT	T3 BERT	Comments
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0-3)	_	Limited algorithms
E1 or T1	Yes (port 0-3)	Yes (port 0-3)	Single port at a timeLimited algorithms
E3 or T3	Yes (port 0-3)	Yes (port 0-3)	Single port at a time
Channelized OC12	_	Yes (channel 0–11)	Single channel at a timeLimited algorithmsNo bit count
Channelized STM1	Yes (channel 0–62)		 Multiple channels Only one algorithm No error insert No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	 Multiple ports and channels Limited algorithms for T1 No error insert for T1 No bit count for T1

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see Channelized IQ and IQE Interfaces Properties.

Starting and Stopping 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.

NOTE: 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

show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)

Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters

IN THIS SECTION

- Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters | 566
- Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm | 567
- Fast Ethernet and Gigabit Ethernet Counters | 570

Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

IN THIS SECTION

- Purpose | **566**
- Action | 566

Purpose

To locate LINK alarm and major counters associated with Fast Ethernet and Gigabit Ethernet interfaces.

Action

Table 168 on page 567 provides links and commands for locating LINK alarm and major counters for Fast Ethernet and Gigabit Ethernet interfaces.

Table 168: Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

Tasks	Command or Action
"Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm" on page 567	show interfaces (fe-fpc/pic/port ge-fpc/pic/port) extensive
"Fast Ethernet and Gigabit Ethernet Counters" on page 570	

SEE ALSO

Ethernet Interfaces User Guide for Routing Devices

Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm

IN THIS SECTION

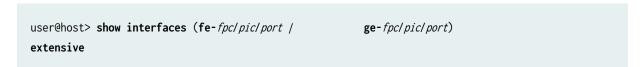
- Problem | 567
- Solution | **567**

Problem

Description

To display the Fast Ethernet or Gigabit Ethernet LINK alarm, use the following Junos OS command-line interface (CLI) operational mode command:

Solution



Sample Output

The following sample output is for a Fast 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
                                                       1602563
   Unicast packets
                                       972137
   Broadcast packets
                                                           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 Fast Ethernet or 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 for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces" on page 517 for information on conducting a loopback test.

NOTE: Since link status is polled once every second, some items that require fast link down detection, such as Multiprotocol Label Switching (MPLS) fast reroute, take longer to execute.

SEE ALSO

Fast Ethernet and Gigabit Ethernet Counters | 570

Ethernet Interfaces User Guide for Routing Devices

Fast Ethernet and Gigabit Ethernet Counters

IN THIS SECTION

- Problem | 570
- Solution | **570**

Problem

Description

Table 169 on page 570 shows the major counters that appear in the output for the show interfaces fefpcl picl port extensive and the show interfaces ge-fpcl picl port extensive commands. These counters generally increment when there is a problem with a Fast Ethernet or Gigabit Ethernet interface. In the Counters column, the counters are listed in the order in which they are displayed in the output.

Solution

Table 169: Major Fast Ethernet and Gigabit Ethernet Counters

Counter	Description	Reason for Increment
Input Errors:		
Errors	The sum of the incoming frame terminates and frame check sequence (FCS) errors.	
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.

Table 169: Major Fast Ethernet and Gigabit Ethernet Counters (Continued)

Counter	Description	Reason for Increment		
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.		
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 terminates and FCS errors.			
Collisions	The number of Ethernet collisions.	The Fast Ethernet PIC supports only full-duplex operation, so this number should always remain 0. If it is incrementing, there is a software bug.		

Number of packets with an

indicated error that the filter

was given to transmit.

Output packet

error count

Table 169: Major Fast Ethernet and Gigabit Ethernet Counters (Continued)			
Counter	Description	Reason for Increment	
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 C	Counters		
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.	

These packets are usually aged packets or are the result of a

the value of this field should not increment.

bandwidth problem on the FPC hardware. On a normal system,

Table 169: Major Fast Ethernet and Gigabit Ethernet Counters (Continued)

Counter	Description	Reason for 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

Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters | 566

Ethernet Interfaces User Guide for Routing Devices

Understanding Interfaces on ACX Series Universal Metro Routers

ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping

Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State

IN THIS SECTION

- Problem | **573**
- Diagnosis | 574
- Resolution | 574

Problem

Description

10-Gigabit Ethernet port is stuck in DPC or PIC down state.

Environment

Juniper Networks T Series and MX Series routers. Refer to the related documentation section for more information.

Symptoms

The device has failed to initialize because the Ethernet port is down.

Diagnosis

Try disabling and reenabling the interface and resetting the transceiver and cable. If the interface remains down, it can be stuck in DPC or PIC down state.

Does the router function normally after disabling and reenabling the interface and resetting the transceiver and cable?

Yes:

The system is not stuck in DPC or PIC down state. Disabling and reenabling the interface or resetting the transceiver, and cable resolved the issue.

No:

The interface might be stuck in DPC or PIC down state. Refer to the "To resolve the issue" on page 574 section for recovery options.

Resolution

To resolve the issue

From the aforementioned diagnosis, you ascertain that the interface is stuck in DPC or PIC down state.

This is not a hardware defect. Implement one of the following solutions on the backup Routing Engine to resolve this issue:

- · 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 (T Series Routers)

```
user@host1# request chassis pic fpc-slot x pic-slot y offline
user@host1# request chassis pic fpc-slot x pic-slot y online
```

4. Reset the PIC (MX Series Routers)

```
user@host1# request chassis fpc slot x offline
user@host1# request chassis fpc slot x online
```

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 2

DPCs Supported on MX240, MX480, and MX960 Routers

T1600 PICs Supported

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test

IN THIS SECTION

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | 576

- Guidelines to perform Bidirectional Diagnostics using Remote Loopback | 580
- Clearing the Interface Statistics | 583

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test

The Pseudo Random Binary Sequence (PRBS) test is a standard feature to verify link quality and transceiver operation. There are two possible diagnostic scenarios:

- Bidirectional verification using remote loopback. Test sequence is initiated from the local end.
 Remote end is configured with the loopback and analysis of the test pattern is performed at local end as well.
- Unidirectional. Test sequence is initiated from the local end. Test pattern is analyzed by the remote end.

In the first case, verification is bidirectional and loopback support is required on the remote end.

The following table lists the entity that enables the PRBS test on various MICs:

MIC Type	PRBS Test Engine	Remarks
JNP10003-LC2103	Packet Forwarding Engine (PFE)	-
JNP-MIC1	Packet Forwarding Engine (PFE)	-
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.

The following table mentions the PRBS test details supported on various interfaces:

Interface Type	Interface Name	Interface Lane Characteristics	Remarks
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.

Guidelines to perform 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.

Following are the steps to collect and view the PRBS statistics:

1. Start a TX (direction 0) by issuing the following command:

[edit]

user@host1> test interface et- θ /1/2 prbs-test-start pattern-type 31 direction 0 flip θ

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
```

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 \theta
```

After executing the command, you can 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. After starting the statistics collection, you can 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 \theta
```

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
```

SEE ALSO

prbs-test-start

prbs-test-stop

show interfaces prbs-stats

clear interfaces statistics

Guidelines to perform Bidirectional Diagnostics using 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
```

NOTE: 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
```

NOTE: There is a change in the direction as 1.

4. After starting the statistics collection, you can 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
```

Interface Card Specific differences

While collecting statistics, the JNP-MIC1-MACSEC and JNP-MIC1 MICs behave differently:

 On JNP-MIC1-MACSEC MIC, if RX is not latching to any PRBS signal, then the "state" in the show interfaces interface-name prbs-stats displays as "Disabled" with Error count as 0, where the JNP-MIC1 MIC displays as failed with MAX error count.

For Example: On JNP-MIC1-MACSEC MIC

```
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

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
```

For Example: On JNP-MIC1 MIC

```
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

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
```

- If any mismatch is encountered between the pattern-type and flip between TX and RX, the maximum
 error counts are observed in case of JNP-MIC1 MIC and state "disabled" in case of JNP-MIC1MACSEC MIC (flip is only supported in JNP-MIC1-MACSEC).
- In case of JNP-MIC1 MIC, if TX is interrupted, the RX displays the state as failed with error counts. Even if TX is started again, RX also must to be restarted to work properly. In case of JNP-MIC1-MACSEC MIC, if TX is interrupted, the RX displays the state as "Disabled" with Error Count 0 (point 1) and if TX is started, RX need not be started again.
- If TX or RX is started consecutively without stopping the earlier run, then difference in the behavior of JNP-MIC1-MACSEC and JNP-MIC1 MICs are observed.
- Decision feedback equalization (DFE) tuning is required on JNP-MIC1 MIC to start a PRBS test. But, on JNP-MIC1-MACSEC MIC, the DFE tuning is not required. If PRBS is started again at TX or RX without stopping the earlier run, there will be errors until DFE tuning is completed again, in JNP-MIC1 MIC. JNP-MIC1-MACSEC MIC does not show this behavior as there is no DFE tuning involved.
- You must disable Forward Error Correction (FEC), if loopback is configured in the router with JNP-MIC1 MIC 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
```

Clearing 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 vales 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
```

```
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.

RELATED DOCUMENTATION

prbs-test-start

prbs-test-stop

show interfaces prbs-stats

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | 575



Configuration Statements and Operational Commands

Unused | 587

Junos CLI Reference Overview | 589

Unused

IN THIS SECTION

- Syntax | 587
- Hierarchy Level | 587
- Description | 587
- Options | 588
- Required Privilege Level | 588
- Release Information | 588

Syntax

unused;

Hierarchy Level

[edit chassis fpc *slot-number* pic *pic-number* port *port-number*]
[edit interface *interface-name*]

Description

(PTX10003-80C and PTX10003-160C routers) While setting port 0 to 400-Gigabit Ethernet mode (using QSFP56-DD-400GBASE-LR8 optics), the total bandwidth (speed x number-of-subports) of port 1 has to be less than 100G and port 2 has to configured as 'unused'. When using port 4 as 400G, port 3 has to be configured with total bandwidth and port 2 has to be configured 'unused'. Similarly, with port 5, 9 using 400G, port 6, 8 respectively has to 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 then 100-Gbps, and the middle port (2 between 0~4 or 7 between 5~9) must be set to unused. To set a particular port to unused port, use unused CLI command.

You cannot use port 1,3,6 and/or 7 of the same logical PIC.

(JNP10K-LC1201 on PTX10008 routers) To disable the specified physical port in a PIC. If you disable a port, no interfaces (channelized or non-channelized) are created for the port, irrespective of the port profile configuration configured on the port.

The edit chassis hierarchy is available in 20.2R1 Release on PTX10001-36MR. From 20.3R1 onwards you must use only the edit interface interface interface hierarchy.

Options

None

Required Privilege Level

interface—To view this statement in the configuration.

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

Release Information

Statement introduced in Junos OS Evolved Release 19.3R1.

Hierarchy introduced in Junos OS Evolved Release 20.1R1 for JNP10K-LC1201 line card on PTX10008 routers.

RELATED DOCUMENTATION

speed (Ethernet)

number-of-sub-ports

PTX10003 Port Panel

Port Speed on PTX10003

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Learn about the syntax and options that make up the statements and commands and 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