

Junos® OS

Ethernet Interfaces User Guide for Routing Devices

Published
2021-04-18

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 © 2021 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 | xxvi

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 | 26

Configuring the MAC Address on the Management Ethernet Interface | 27

Power over Ethernet (PoE) on ACX Series Routers | 29

Understanding PoE on ACX Series Universal Metro Routers | 29

Example: Configuring PoE on ACX2000 Routers | 33

Requirements | 33

Overview | 33

Configuration | 34

Verification | 37

Example: Disabling a PoE Interface on ACX2000 Routers | 39

Requirements | 40

Overview | 40

Configuration | 40

Verification | 41

Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers | 41

Point-to-Point Protocol over Ethernet (PPPoE) | 43

PPPoE Overview | 43

Configuring PPPoE | 48

Overview | 48

Setting the Appropriate Encapsulation on the PPPoE Interface | 49

Configuring PPPoE Encapsulation on an Ethernet Interface | 50

Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface | 50

Configuring the PPPoE Underlying Interface | 51

Identifying the Access Concentrator | 51

Configuring the PPPoE Automatic Reconnect Wait Timer | 52

Configuring the PPPoE Service Name | 52

Configuring the PPPoE Server Mode | 52

Configuring the PPPoE Client Mode | 53

Configuring the PPPoE Source and Destination Addresses | 53

Deriving the PPPoE Source Address from a Specified Interface | 53

Configuring the PPPoE IP Address by Negotiation | 54

Configuring the Protocol MTU PPPoE | 54

Example: Configuring a PPPoE Server Interface on an M120 or M320 Router | 55

Disabling the Sending of PPPoE Keepalive Messages | 55

Verifying a PPPoE Configuration | 56

Tracing PPPoE Operations | 57

Configuring the PPPoE Trace Log Filename | 58

Configuring the Number and Size of PPPoE Log Files | 58

- Configuring Access to the PPPoE Log File | 59
- Configuring a Regular Expression for PPPoE Lines to Be Logged | 59
- Configuring the PPPoE Tracing Flags | 59

Configuring Aggregated Ethernet Interfaces | 61

Aggregated Ethernet Interfaces | 61

- What Are Aggregated Ethernet Interfaces? | 62
- Configuration Guidelines for Aggregated Ethernet Interfaces | 63
- Configure Aggregated Ethernet Interfaces | 64
- Mixed-Mode and Mixed-Rate Aggregated Ethernet Interfaces | 67
- Platform Support for Mixed Aggregated Ethernet Bundles | 67
- Configuration Guidelines for Mixed-Rate Aggregated Ethernet Links | 70
- Configure Mixed-Rate Aggregated Ethernet Interfaces | 71
- What Is Link Aggregation Control Protocol? | 72
- Configuration Guidelines for LACP | 73
- Configure LACP | 74
- Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 76
- Example: Configure Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links | 77
 - Requirements | 81
 - Overview | 81
 - Verification | 82
- Independent Micro-BFD Sessions for LAG | 89
- Configuration Guidelines for Micro-BFD Sessions | 90
- Example: Configure Independent Micro-BFD Sessions for LAG | 91
 - Requirements | 91
 - Overview | 92
 - Configuration | 92
 - Verification | 100
- MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet Interfaces | 104
- What Is Enhanced LAG? | 105

Link Protection of Aggregated Ethernet Interfaces | 106

- Configuring Aggregated Ethernet Link Protection | 106
 - Configuring Link Protection for Aggregated Ethernet Interfaces | 107
 - Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces | 108

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link | 108

Disabling Link Protection for Aggregated Ethernet Interfaces | 109

Configuring Aggregated Ethernet Minimum Links | 109

Example: Configuring Aggregated Ethernet Link Protection | 110

Scheduling on Aggregated Ethernet Interfaces | 111

Configuring Shared Scheduling on Aggregated Ethernet Interfaces | 111

Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection | 112

Load Balancing on Aggregated Ethernet Interfaces | 113

Load Balancing and Ethernet Link Aggregation Overview | 114

Understanding Aggregated Ethernet Load Balancing | 115

Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data | 117

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | 120

Configuring Adaptive Load Balancing | 121

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers | 123

Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview | 123

Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers | 123

Configuring Symmetrical Load Balancing on Trio-Based MPCs | 127

Example Configurations | 129

Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers | 130

Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers | 133

Example: Configuring Aggregated Ethernet Load Balancing | 136

Example: Configuring Aggregated Ethernet Load Balancing | 136

Performance Monitoring on Aggregated Ethernet Interfaces | 155

ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces Overview | 155

Guidelines for Configuring Performance Monitoring Functionalities on Aggregated Ethernet Interfaces | 158

Periodic Packet Management | 159

Understanding Periodic Packet Management on MX Series Routers | 160

Configuring Periodic Packet Management on MX Series Routers | 160

Identifying Periodic Packet Management Mode | 161

Enabling Centralized Periodic Packet Management | 162

Understanding Ethernet Link Aggregation on ACX Series Routers | 163

2

Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Interfaces | 173

Configuring 10-Gigabit Ethernet PICs | 173

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 174

12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 178

24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 181

P2-10G-40G-QSFPP PIC Overview | 182

Configuring the P2-10G-40G-QSFPP PIC | 193

Configuring the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 193

Configuring the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 194

Configuring the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 194

Configuring the PIC at Port Group Level | 195

Configuring Framing Mode on P2-10G-40G-QSFPP PIC | 196

Example: Configuring the P2-10G-40G-QSFPP PIC | 197

Requirements | 198

Overview | 198

Configuration | 198

Framing Overview | 202

Understanding WAN Framing | 203

Configuring Ethernet Framing | 203

Modes of Operation | 204

Configuring Mixed-Rate Mode Operation | 205

Configuring Line-Rate Mode on PICs Supporting Oversubscription | 207

Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC | 207

Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 209

Gigabit Ethernet Notification of Link Down Alarm Overview | 212

Notification of Link Down for Optics Options Overview | 213

Configuring Gigabit Ethernet Notification of Link Down Alarm | 213

Configuring Link Down Notification for Optics Options Alarm or Warning | 214

Configuring 40-Gigabit Ethernet PICs | 215

40-Gigabit Ethernet PIC Overview | 215

Configuring 40-Gigabit Ethernet PICs | 217

Configuring 100-Gigabit Ethernet MICs/PICs | 218

- 100-Gigabit Ethernet Interfaces Overview | **219**
- MPC3E MIC Overview | **222**
- 100-Gigabit Ethernet Type 4 PIC with CFP Overview | **224**
- Configuring 100-Gigabit Ethernet Type 4 PIC With CFP | **227**
- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | **231**
- 100-Gigabit Ethernet Type 5 PIC with CFP Overview | **233**
- 100-Gigabit Ethernet Interfaces Interoperability | **235**
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | **238**
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | **239**
 - Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP | **239**
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | **241**
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | **242**
 - Configuring SA Multicast Bit Steering Mode on 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP | **243**
 - Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface | **244**
- Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | **246**
 - Transporting Legacy Traffic over Packet Switched Networks | **246**
 - Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH Traffic | **249**
 - Requirements for Configuration of the Smart-SFPs on MX Series Routers | **249**
 - Overview | **250**
 - Configuring the DS3 Smart SFP | **250**
 - Verification | **252**
 - Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy SDH Traffic | **254**
 - Requirements for Configuration of the Smart-SFPs on MX Series Routers | **255**
 - Overview | **255**
 - Configuring the STM1 Smart SFP | **255**
 - Verification | **257**
- Configuring Layer 2 Overhead Attribute in Interface Statistics | **259**
 - Accounting of the Layer 2 Overhead Attribute in Interface Statistics | **260**

Configuring Layer 2 Overhead Accounting in Interface Statistics | 263

 | Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level | 263

Verifying the Accounting of Layer 2 Overhead in Interface Statistics | 264

Configuring Gigabit Ethernet Policers | 268

Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 268

Configuring Gigabit Ethernet Policers | 271

 | Overview | 271

 | Configuring a Policer | 272

 | Specifying an Input Priority Map | 273

 | Specifying an Output Priority Map | 273

 | Applying a Policer | 274

 | Configuring MAC Address Filtering | 276

 | Example: Configuring Gigabit Ethernet Policers | 277

Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 279

 | Overview | 279

 | Configuring a Policer | 280

 | Applying a Policer | 281

 | Example: Configuring and Applying a Policer | 282

Gigabit Ethernet Autonegotiation | 284

Gigabit Ethernet Autonegotiation Overview | 284

Configuring Gigabit Ethernet Autonegotiation | 284

 | Configuring Gigabit Ethernet Autonegotiation with Remote Fault | 285

 | Configuring Flow Control | 285

 | Configuring Autonegotiation Speed on MX Series Routers | 285

 | Displaying Autonegotiation Status | 286

Configure Port Speed | 293

Port Speed | 293

 | Port Speed Overview | 294

 | Port Speed on MX10003 MPC Overview | 306

 | Port Speed on MPC10E-10C-MRATE Overview | 312

 | Port Speed on MPC10E-15C-MRATE Overview | 317

 | Port Speed on MX2K-MPC11E Overview | 322

 | Port Speed on PTX10K-LC1201 Overview | 334

 | Port Speed on ACX7100-48L Router Overview | 345

Interface Naming Conventions for ACX7100-48L | **347**

Port Speed on ACX710 Router Overview | **348**

Port Speed on PTX10001-36MR Router Overview | **350**

Port Speed on PTX10K-LC1202-36MR Overview | **361**

Introduction to Port Speed | **370**

Port Speed for Routing Devices | **370**

Guidelines for Configuring Port Speed on Routing Devices | **374**

MX10003 MPC Port Speed Overview | **376**

MX204 Router Port Speed Overview | **382**

PTX10003 Router Port Speed Overview | **390**

PTX10K-LC1201 Port Speed Overview | **392**

Interface Naming Conventions | **394**

Interface Naming Conventions for MPC7E-MRATE | **394**

Interface Naming Conventions for MIC-MRATE | **396**

Interface Naming Conventions for MX10003 MPC | **400**

Interface Naming Conventions for PTX10K-LC2101 | **403**

Interface Naming Conventions for MIC-MACSEC-20GE | **405**

Interface Naming Conventions for PTX10K-LC1201 | **407**

Interface Mapping and Modulation Format for ACX5448-D | **416**

Configuring Port Speed on Routing Devices and Line Cards | **418**

Configuring Port Speed | **418**

Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers | **419**

Configuring Port Speed on MIC-MRATE to Enable Different Port Speeds | **423**

Configuring Port Speed on MPC7E (Multi-Rate) to Enable Different Port Speeds | **427**

Configuring Port Speed on MX10003 MPC to Enable Different Port Speeds | **432**

Configuring Port Speed on MX204 to Enable Different Port Speeds | **436**

Configuring Port Speed on PTX10K-2101 MPC to Enable Different Port Speeds | **440**

Configure Port Speed on ACX5448-D and ACX5448-M Routers | **444**

Channelize Interfaces on ACX5448-D and ACX5448-M Routers | **445**

Configuring Port Speed on PTX10K-1201 line card to Enable Different Port Speeds | **448**

Configuring the Port Speed on the JNP10K-LC1201 by Using New Port Profile Configuration | **452**

Understanding New Port Profile Configuration | **452**

Configuring Port Speed on JNP10K-LC1201 by using New Port Profile Configuration | **454**

Configuring Optical Transport network | 458

Understanding Optical Transport Network (OTN) | **458**

10-Gigabit Ethernet OTN Options Configuration Overview	459
100-Gigabit Ethernet OTN Options Configuration Overview	459
Understanding the P1-PTX-24-10G-W-SFPP PIC	461
Understanding the features of ACX6360	466
Understanding the P2-100GE-OTN PIC	469
Understanding the MIC3-100G-DWDM MIC	474
Understanding the PTX-5-100G-WDM PIC	478
Understanding the PTX10K-LC1104 Line Card	482
Interface Mapping and Modulation format for ACX6360	485
Interface Mapping and Modulation format for PTX10K-LC1104 Line Card	488

Supported OTN and Optics Options | 491

Supported OTN Options on PTX Series Routers	491
Supported OTN Options on MX Series Routers	503
Supported OTN Options on ACX6360 Routers	515
Supported OTN Options on ACX5448-D Routers	519
Supported OTN Options on PTX10008 and PTX10016 Series Routers	524
Supported Optics Options on ACX6360 and ACX5448-D Routers	533
Supported Optics Options on PTX10008 and PTX10016 Series Routers	537

Forward Error Correction (FEC) and Bit Error Rate (BER) | 543

Understanding Pre-FEC BER Monitoring and BER Thresholds	543
Supported Forward Error Correction Modes on MX Series Routers	549
Supported Forward Error Correction Modes on PTX Series Routers	549
Supported Forward Error Correction Modes on ACX6360 Router	550
Supported FEC Modes on ACX5448-D Router	551

Dense Wavelength Division Multiplexing (DWDM) Interface Wavelength | 552

Ethernet DWDM Interface Wavelength Overview	552
Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength	553

Configuring OTN | 556

Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC	557
Configuring OTN Interfaces on P1-PTX-2-100G-WDM	561
Configuring OTN Interfaces on P2-100GE-OTN PIC	563
Configuring OTN Interfaces on MIC3-100G-DWDM MIC	569
Configuring OTN Interfaces on PTX-5-100G-WDM PIC	575
Configuring OTN Interface Options on PTX10K-LC1104	581

ODU Path Delay Measurement for Performance Monitoring | 587

Understanding ODU Path Delay Measurement | 587

Enabling ODU Path Delay Measurement | 589

Disabling ODU Path Delay Measurement | 591

Troubleshooting Information

Monitoring and Troubleshooting Ethernet Interfaces | 595

Passive Monitoring on Ethernet Interfaces Overview | 595

Enabling Passive Monitoring on Ethernet Interfaces | 597

Link Degrade Monitoring Overview | 600

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 604

Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 605

Monitor Fast Ethernet and Gigabit Ethernet Interfaces | 606

Display the Status of Fast Ethernet Interfaces | 607

Display the Status of Gigabit Ethernet Interfaces | 609

Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface | 610

Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface | 612

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 617

Fiber-Optic Ethernet Interface Specifications | 619

Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 621

Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 621

Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface | 623

Create a Loopback | 624

Create a Physical Loopback for a Fiber-Optic Interface | 624

Create a Loopback Plug for an RJ-45 Ethernet Interface | 625

Configure a Local Loopback | 626

Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 627

Configure a Static Address Resolution Protocol Table Entry | 632

Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 637

Ping the Fast Ethernet or Gigabit Ethernet Interface | 637

Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 639

Diagnose a Suspected Circuit Problem | 642

Performing Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 643

Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | **643**

Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | **645**

Create a Loopback | **646**

 Create a Physical Loopback for a Fiber-Optic Interface | **646**

 Configure a Local Loopback | **647**

Verify That the Gigabit Ethernet Interface Is Up | **648**

Configure a Static Address Resolution Protocol Table Entry | **654**

Clear Gigabit Ethernet Interface Statistics | **655**

Ping the Gigabit Ethernet Interface | **656**

Check for Gigabit Ethernet Interface Error Statistics | **658**

Diagnose a Suspected Circuit Problem | **662**

Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | **663**

 Configuring Loopback Testing | **663**

 Configuring BERT Testing | **666**

 Starting and Stopping a BERT Test | **670**

Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters | **671**

 Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters | **671**

 Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm | **672**

 Fast Ethernet and Gigabit Ethernet Counters | **675**

Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State | **679**

Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | **681**

 Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | **681**

 Guidelines to perform Bidirectional Diagnostics using Remote Loopback | **686**

 Clearing the Interface Statistics | **689**

4

Configuration Statements and Operational Commands

Configuration Statements (OTN) | **693**

alarm (optics-options) | **694**

backward-frr-enable | **696**

ber-threshold-clear | **698**

ber-threshold-signal-degrade | **701**

bypass | **704**

bytes (otn-options) | 706

cfp-to-et | 707

fec | 709

fec | 712

fec (gigether) | 713

fixed-stuff-bytes | 715

high-polarization | 717

interval | 718

is-ma | 721

laser-enable | 722

line-loopback | 724

local-loopback | 725

monitor-end-point | 727

no-odu-backward-frr-enable | 728

no-odu-signal-degrade-monitor-enable | 729

number-of-frames | 731

oc192 | 732

odu-delay-management | 733

odu-backward-frr-enable | 735

odu-signal-degrade | 736

odu-signal-degrade-monitor-enable | 738

odu-ttim-action-enable | 739

otu-ttim-action-enable | 741

otu4 | 742

pass-through | 744

prbs | 745

preemptive-fast-reroute | 747

rate | 748

remote-loop-enable | 750

signal-degrade | 752

signal-degrade-monitor-enable | 753

start-measurement | 755

tca | 757

transport-monitoring | 761

trigger | 762

tti | 767

tx-power | 768

warning | 770

wavelength | 771

Configuration Statements | 777

accept-source-mac | 784

access-concentrator | 786

account-layer2-overhead (PIC Level) | 788

adaptive | 789

address | 791

age | 795

agent-specifier | 796

aggregate (Gigabit Ethernet CoS Policer) | 798

aggregated-ether-options | 800

alarms | 802

allow-remote-loopback | 803

asynchronous-notification | 804

auto-negotiation | **806**

auto-reconnect | **808**

bandwidth-limit (Policer for Gigabit Ethernet Interfaces) | **810**

bert-algorithm | **811**

bert-error-rate | **814**

bert-period | **816**

bridge-domain | **818**

burst-size-limit (Policer for Gigabit Ethernet Interfaces) | **820**

centralized | **821**

ces-psn-channel (tdm-options) | **823**

cfp-to-et | **824**

classifier | **826**

client | **827**

community-vlans (MX Series) | **828**

connectivity-fault-management | **830**

delay (PPPoE Service Name Tables) | **834**

destination (IPCP) | **835**

device-count | **837**

drop (PPPoE Service Name Tables) | **838**

dynamic-profile (PPPoE Service Name Tables) | **840**

egress-policer-overhead | **842**

encapsulation (Logical Interface) | **843**

encapsulation | **848**

ether-options | **856**

ethernet (Chassis) | **865**

ethernet-policer-profile | **867**

evcs | **869**

family | **870**

family | **872**

fastether-options | **878**

flow-control | **880**

fnp | **882**

force-up | **883**

forwarding-class (Gigabit Ethernet IQ Classifier) | **885**

forwarding-mode (100-Gigabit Ethernet) | **886**

forwarding-mode (PTX Series Packet Transport Routers) | **888**

frame-error | **889**

frame-period | **891**

frame-period-summary | **892**

framing (10-Gigabit Ethernet Interfaces) | **894**

gether-options | **896**

hash-key (Chassis LAG) | **905**

hold-time up | **907**

ieee802.1p | **908**

ignore-l3-incompletes | **910**

inet (chassis) | **911**

ingress-policer-overhead | **912**

ingress-rate-limit | **915**

inline | **916**

input-policer | **918**

input-priority-map | **920**

input-three-color | **921**

input-vlan-map (Aggregated Ethernet) | 923

interface (LLDP) | 924

interface (OAM Link-Fault Management) | 927

interface-group | 929

interface-group-down | 930

interface-none | 931

isolated-vlan (MX Series) | 932

iwf-params (tdm-options) | 934

lACP (802.3ad) | 936

lACP (Aggregated Ethernet) | 937

lACP | 940

lACP (Protocols) | 942

lane | 944

lane-all | 947

layer2-policer | 950

link-adjacency-loss | 952

link-discovery | 954

link-degrade-monitor | 955

link-down | 958

link-event-rate | 959

link-fault-management | 960

link-mode | 962

link-protection | 964

link-protection (non-LACP) | 966

link-protection (Protocols LACP) | 968

link-speed (Aggregated Ethernet) | 969

link-speed (Aggregated SONET/SDH) | **972**

lmi (Ethernet OAM) | **973**

load-balance | **975**

load-balance-stateful (Aggregated Ethernet Interfaces) | **977**

load-type (Aggregated Ethernet Interfaces) | **979**

local-bias (ae load-balance) | **981**

logical-tunnel-options | **982**

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | **984**

loopback (Local and Remote) | **986**

loopback-tracking | **988**

loss-priority | **989**

loopback-remote | **990**

mac | **991**

mac-address (Accept Source Mac) | **993**

mac-learn-enable | **994**

mac-validate | **996**

master-only | **997**

max-sessions (PPPoE Service Name Tables) | **998**

max-sessions-vs-a-ignore (Static and Dynamic Subscribers) | **1000**

maximum-links | **1002**

mc-ae | **1004**

minimum-bandwidth (aggregated Ethernet) | **1010**

minimum-links | **1012**

mixed-rate-mode | **1014**

mtu | **1016**

mru | **1020**

multicast-statistics | **1021**

multiservice | **1023**

negotiate-address | **1025**

negotiation-options | **1026**

no-adaptive | **1027**

no-allow-link-events | **1029**

no-auto-mdix | **1030**

no-keepalives | **1031**

no-pre-classifier | **1033**

no-send-pads-ac-info | **1034**

no-send-pads-error | **1036**

non-revertive (Chassis) | **1037**

non-revertive (Interfaces) | **1038**

number-of-ports | **1040**

number-of-sub-ports | **1042**

oam | **1044**

optics-options | **1048**

otn-options | **1051**

output-policer | **1054**

output-priority-map | **1055**

output-three-color | **1057**

output-vlan-map (Aggregated Ethernet) | **1058**

pado-advertise | **1060**

passive-monitor-mode | **1061**

payload | **1062**

pdu-interval | **1064**

pdu-threshold | **1066**

per-flow (Aggregated Ethernet Interfaces) | **1067**

periodic | **1069**

pic-mode | **1071**

policer (CFM Firewall) | **1073**

policer (CoS) | **1074**

policer (MAC) | **1076**

port-priority | **1078**

pp0 (Dynamic PPPoE) | **1079**

ppm (Ethernet Switching) | **1082**

pppoe-options | **1083**

pppoe-underlying-options (Static and Dynamic Subscribers) | **1085**

preferred-source-address | **1087**

premium (Output Priority Map) | **1089**

premium (Policer) | **1090**

protocol-down | **1091**

premium (Output Priority Map) | **1092**

premium (Policer) | **1094**

proxy | **1095**

rebalance (Aggregated Ethernet Interfaces) | **1097**

receive-options-packets | **1098**

receive-ttl-exceeded | **1100**

recovery | **1101**

remote-loopback | **1103**

restore-interval | **1104**

revertive | **1106**

routing-instance | **1107**

routing-instance (PPPoE Service Name Tables) | **1109**

rx-enable | **1110**

rx-max-duration | **1112**

sa-multicast (100-Gigabit Ethernet) | **1113**

sa-multicast (PTX Series Packet Transport Routers) | **1115**

send-critical-event | **1116**

server | **1117**

service (PPPoE) | **1119**

service-name | **1121**

service-name-table | **1122**

service-name-tables | **1124**

session-expiry (MX Series in Enhanced LAN Mode) | **1126**

sonet | **1128**

source-address-filter | **1129**

source-filtering | **1131**

speed | **1132**

speed | **1134**

speed (Ethernet) | **1137**

speed (MX Series DPC) | **1148**

speed (Gigabit Ethernet interface) | **1150**

speed (24-port and 12-port 10 Gigabit Ethernet PIC) | **1152**

static-interface | **1154**

switch-options | **1156**

switch-port | **1157**

symbol-period | **1158**

symmetric-hash | **1160**

sync-reset | **1161**

syslog (OAM Action) | **1163**

system-id | **1164**

system-priority | **1166**

system-priority | **1167**

targeted-options (Grouping Subscribers by Bandwidth Usage) | **1169**

targeted-options (Manual Targeting) | **1171**

targeted-distribution | **1173**

targeted-options | **1174**

tdm-options (Interfaces) | **1176**

terminate (PPPoE Service Name Tables) | **1178**

thresholds | **1179**

traceoptions (LLDP) | **1181**

traceoptions (Individual Interfaces) | **1185**

traceoptions (LACP) | **1194**

traceoptions (PPPoE) | **1197**

tx-duration | **1200**

tx-enable | **1202**

underlying-interface | **1203**

unit | **1205**

unnumbered-address (Dynamic Profiles) | **1215**

unnumbered-address (PPP) | **1219**

Unused | **1220**

virtual-control-channel | **1222**

virtual-switch | **1224**

vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) | 1225

vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) | 1226

Operational Commands | 1228

clear interfaces interface-set statistics | 1229

clear interfaces interval | 1231

clear interfaces aeX forwarding-options load-balance state | 1233

clear interfaces aggregate forwarding-options load-balance state | 1235

clear interfaces transport pm | 1236

clear protection-group ethernet-ring statistics | 1238

prbs-test-start | 1239

prbs-test-stop | 1242

request interface link-degrade-recover | 1243

request interface mc-ae switchover (Multichassis Link Aggregation) | 1247

request interface (revert | switchover) (Aggregated Ethernet Link Protection) | 1250

request lacp link-switchover | 1251

show chassis hardware | 1253

show chassis pic | 1274

show ethernet-switching redundancy-groups | 1311

show interfaces (Adaptive Services) | 1317

show interfaces (Aggregated Ethernet) | 1328

show interfaces demux0 (Demux Interfaces) | 1352

show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) | 1368

show interfaces (far-end-interval) | 1432

show interfaces (Fast Ethernet) | 1434

show interfaces | 1463

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) | **1602**

show interfaces (PPPoE) | **1627**

show interfaces interface-set (Ethernet Interface Set) | **1645**

show interfaces interface-set queue | **1651**

show interfaces interval | **1663**

show interfaces irb | **1667**

show interfaces mac-database | **1680**

show interfaces mc-ae | **1688**

show interfaces port-profile | **1695**

show interfaces prbs-stats | **1703**

show interfaces smart-sfp-defects | **1705**

show interfaces smart-sfp-statistics | **1719**

show interfaces transport pm | **1737**

show l2-learning instance | **1747**

show l2-learning redundancy-groups | **1750**

show lacp interfaces | **1756**

show pppoe interfaces | **1765**

show pppoe service-name-tables | **1771**

show pppoe sessions | **1776**

show pppoe statistics | **1778**

show pppoe underlying-interfaces | **1782**

show pppoe version | **1793**

traceroute ethernet | **1795**

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](#)

1

PART

Ethernet Interfaces

[Configuring Ethernet Interfaces | 2](#)

[Configuring Aggregated Ethernet Interfaces | 61](#)

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 | 29](#)
- [Point-to-Point Protocol over Ethernet \(PPPoE\) | 43](#)

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-*fpc/pic/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 **gether-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 ggether-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 **gether-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 ggether-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 *logical-unit-number*].

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

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*

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 **10m** maximum 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

| *speed*

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.nnnn or *nr:nr:nr:nr:nr:nr* 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 **ixgbe0** and **ixgbe1**. For T1600 routers configured in a routing matrix, the internal Ethernet interfaces are **bcm0** and **em1**. 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           up    up   inet    10.34.0.4/8
                                     162.0.0.4/2
                                     inet6   fe80::200:ff:fe22:4/64
                                     fec0::a:22:0:4/64
                                     tnp     0x22000004

```

```
user@host-sfc-0> show interfaces ixgbe1 terse
```

```

Interface          Admin Link Proto   Local                               Remote
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
                                     tnp     0x22000004

```

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.

```
[edit]
user@host# set interfaces fe-5/2/1 unit 0 family inet address local-address
user@host# show
interfaces {
  fe-5/2/1 {
    unit 0 {
      family inet {
        address local-address;
      }
    }
  }
}
```

SEE ALSO

[Management Ethernet Interfaces](#)

[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](#)

Display the Status of Gigabit Ethernet Interfaces

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

Release History Table

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 10m maximum 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](#) | 880

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* ggether-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

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* ggether-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* ggether-options]

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

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*.

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*.

NOTE: On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement at the [edit interfaces *ge-fpc/pic/port* ggether-options] hierarchy level and the **accept-source-mac** statement at the [edit interfaces *ge-fpc/pic/port* ggether-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* ggether-options] hierarchy level and the **accept-source-mac** statement at the [edit interfaces *ge-fpc/pic/port* ggether-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*.

RELATED DOCUMENTATION

[source-address-filter](#)

[Aggregated Ethernet Interfaces](#)

[Ethernet Interfaces Overview](#)

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 gigheter-options]` hierarchy level:

```
[edit interfaces]
et-x/y/z {
  gigheter-options {
    source-filtering;
    source-address-filter {
      mac-address;
    }
  }
}
```

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: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: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

Configuring Gigabit Ethernet Policers

Configuring Gigabit Ethernet Two-Color and Tricolor Policers

Configuring a Policer Overhead

RELATED DOCUMENTATION

[Aggregated Ethernet Interfaces | 61](#)

[Configuring Gigabit Ethernet Policers | 268](#)

Management Ethernet Interfaces

IN THIS SECTION

- [Management Ethernet Interface Overview | 25](#)
- [Configuring a Consistent Management IP Address | 26](#)
- [Configuring the MAC Address on the Management Ethernet Interface | 27](#)

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**.

SEE ALSO

Ethernet Interfaces Overview

Initial Configuration of Ethernet Interfaces

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 rel 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](#)

[Initial Configuration of Ethernet Interfaces](#)

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, **0011.2233.4455**) or *nn:nn:nn:nn:nn:nn* (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

Initial Configuration of Ethernet Interfaces

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 | 29](#)
- [Example: Configuring PoE on ACX2000 Routers | 33](#)
- [Example: Disabling a PoE Interface on ACX2000 Routers | 39](#)
- [Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers | 41](#)

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 | 30](#)
- [PoE Classes and Power Ratings | 30](#)
- [PoE Options | 31](#)

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 1 on page 30 lists the PoE specifications for the ACX2000 routers.

Table 1: PoE Specifications for the ACX2000 Routers

Specifications	For ACX2000 Universal Metro Routers
Supported standards	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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 2 on page 31 lists the classes and their power ratings as specified by the IEEE standards.

Table 2: 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 3 on page 31](#) shows the PoE configuration options and their default settings for the PoE controller and for the PoE interfaces.

Table 3: PoE Configuration Options and Default Settings

Option	Default	Description
PoE Controller Options		
guard-band	0 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 3: PoE Configuration Options and Default Settings (*Continued*)

Option	Default	Description
management	static	<p>Sets the PoE power management mode for the router. The power management mode determines how power to a PoE interface is allocated:</p> <ul style="list-style-type: none"> • 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.
telemetries	Not included in default configuration	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

[Power over Ethernet \(PoE\) User Guide for EX Series Switches](#)

Example: Configuring PoE on ACX2000 Routers

IN THIS SECTION

- [Requirements | 33](#)
- [Overview | 33](#)
- [Configuration | 34](#)
- [Verification | 37](#)

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 4 on page 34](#) details the topology used in this configuration example.

Table 4: 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 | 35
- Results | 37

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.

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

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 | 37](#)
- [Verifying the Telemetry Data \(History\) for the Specified Interface | 38](#)
- [Verifying PoE Global Parameters | 39](#)

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.

```
user@host> show poe controller
Controller  Maximum   Power      Guard   Management  Status   Lldp
index      power     consumption band     Management  Status   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*

Example: Disabling a PoE Interface on ACX2000 Routers

IN THIS SECTION

- [Requirements | 40](#)
- [Overview | 40](#)
- [Configuration | 40](#)
- [Verification | 41](#)

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*.

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 | 40](#)

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 interface
```

Interface	Admin status	Oper status	Max power	Priority	Power consumption	Class
ge-0/1/3	Disabled	Disabled	32.0W	Low	0.0W	0
ge-0/1/7	Disabled	Disabled	32.0W	Low	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 | 42](#)
- [Solution | 42](#)

Problem

Description

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

Solution

Check for the items shown in [Table 5 on page 42](#).

Table 5: 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 <code>show poe interface</code> 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 <code>show poe interface</code> 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 <code>show poe telemetries</code> 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

Troubleshooting PoE Interfaces

Point-to-Point Protocol over Ethernet (PPPoE)

IN THIS SECTION

- [PPPoE Overview | 43](#)
- [Configuring PPPoE | 48](#)
- [Disabling the Sending of PPPoE Keepalive Messages | 55](#)
- [Verifying a PPPoE Configuration | 56](#)
- [Tracing PPPoE Operations | 57](#)

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 | 44](#)
- [PPPoE Stages | 45](#)
- [Optional CHAP Authentication | 47](#)

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

PPPoE Session Stage

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

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

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

Optional CHAP Authentication

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

If you configure an interface to handle incoming CHAP packets only (by including the **passive** statement at the **[edit interfaces *interface-name* ppp-options chap]** hierarchy level), the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not include the **passive** statement, the interface always challenges its peer.

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*](#)

[*Disabling the Sending of PPPoE Keepalive Messages*](#)

[*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

Troubleshooting PPPoE Service Name Tables

Verifying a PPPoE Configuration

[Ethernet Interfaces User Guide for Routing Devices](#)

Configuring PPPoE

IN THIS SECTION

- [Overview | 48](#)
- [Setting the Appropriate Encapsulation on the PPPoE Interface | 49](#)
- [Configuring PPPoE Encapsulation on an Ethernet Interface | 50](#)
- [Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface | 50](#)
- [Configuring the PPPoE Underlying Interface | 51](#)
- [Identifying the Access Concentrator | 51](#)
- [Configuring the PPPoE Automatic Reconnect Wait Timer | 52](#)
- [Configuring the PPPoE Service Name | 52](#)
- [Configuring the PPPoE Server Mode | 52](#)
- [Configuring the PPPoE Client Mode | 53](#)
- [Configuring the PPPoE Source and Destination Addresses | 53](#)
- [Deriving the PPPoE Source Address from a Specified Interface | 53](#)
- [Configuring the PPPoE IP Address by Negotiation | 54](#)
- [Configuring the Protocol MTU PPPoE | 54](#)
- [Example: Configuring a PPPoE Server Interface on an M120 or M320 Router | 55](#)

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](#)" and "[Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface](#)".

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*]
- [edit logical-systems *logical-system-name* interfaces pp0 unit *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, **lo0.0**—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-1/0/0.0;
      server;
    }
    ppp-options {
    }
    family inet {
      address 22.2.2.1/32 {
        destination 22.2.2.2;
      }
    }
  }
}
```

RELATED DOCUMENTATION

[Ethernet Interfaces Overview](#)

[Initial Configuration of Ethernet Interfaces](#)

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](#)

[Ethernet Interfaces User Guide for Routing Devices](#)

Verifying a PPPoE Configuration

IN THIS SECTION

- [Purpose | 56](#)
- [Action | 56](#)

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

[Ethernet Interfaces User Guide for Routing Devices](#)

Tracing PPPoE Operations

IN THIS SECTION

- [Configuring the PPPoE Trace Log Filename | 58](#)
- [Configuring the Number and Size of PPPoE Log Files | 58](#)
- [Configuring Access to the PPPoE Log File | 59](#)
- [Configuring a Regular Expression for PPPoE Lines to Be Logged | 59](#)
- [Configuring the PPPoE Tracing Flags | 59](#)

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
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
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
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
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](#)

**RELATED DOCUMENTATION**

| [PPPoE Overview](#)

**Release History Table**

| 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 logical-unit-number] hierarchy level or at the [edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number] 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.                                                                                                                                                                         |

# Configuring Aggregated Ethernet Interfaces

## IN THIS CHAPTER

- [Aggregated Ethernet Interfaces | 61](#)
- [Link Protection of Aggregated Ethernet Interfaces | 106](#)
- [Scheduling on Aggregated Ethernet Interfaces | 111](#)
- [Load Balancing on Aggregated Ethernet Interfaces | 113](#)
- [Performance Monitoring on Aggregated Ethernet Interfaces | 155](#)
- [Periodic Packet Management | 159](#)
- [Understanding Ethernet Link Aggregation on ACX Series Routers | 163](#)

## 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? | 62](#)
- [Configuration Guidelines for Aggregated Ethernet Interfaces | 63](#)
- [Configure Aggregated Ethernet Interfaces | 64](#)
- [Mixed-Mode and Mixed-Rate Aggregated Ethernet Interfaces | 67](#)
- [Platform Support for Mixed Aggregated Ethernet Bundles | 67](#)
- [Configuration Guidelines for Mixed-Rate Aggregated Ethernet Links | 70](#)

- [Configure Mixed-Rate Aggregated Ethernet Interfaces | 71](#)
- [What Is Link Aggregation Control Protocol? | 72](#)
- [Configuration Guidelines for LACP | 73](#)
- [Configure LACP | 74](#)
- [Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 76](#)
- [Example: Configure Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links | 77](#)
- [Independent Micro-BFD Sessions for LAG | 89](#)
- [Configuration Guidelines for Micro-BFD Sessions | 90](#)
- [Example: Configure Independent Micro-BFD Sessions for LAG | 91](#)
- [MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet Interfaces | 104](#)
- [What Is Enhanced LAG? | 105](#)

## What Are Aggregated Ethernet Interfaces?

### IN THIS SECTION

- [Benefits | 63](#)

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 | 63](#)

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 6 on page 64](#) 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 6: 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, MX960, MX10003, MX10008, MX10016, MX2010, and MX2020 | 64                         | 128 (Before 14.2R1)<br>1000 (14.2R1 and later)  |

[Table 7 on page 64](#) 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 7: 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               |

## Configure Aggregated Ethernet Interfaces

[Table 8 on page 65](#) describes the steps to configure aggregated Ethernet interfaces on your routing device.

**Table 8: Aggregated Ethernet Interfaces Configuration**

| Configuration Step                                                                                                                                                                                                                                                                                                                                                  | Command                                                                                                     |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| <p>Step 1: Specify the number of aggregated Ethernet bundles you want on your device. If you specify the <b>device-count</b> value as 2, you can configure two aggregated bundles.</p>                                                                                                                                                                              | <pre>[edit chassis aggregated-devices ethernet] user@host# <b>set device-count <i>number</i></b></pre>      |
| <p>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.</p>                                                                                                                                                                                 | <pre>[edit interfaces ] user@host# <b>set <i>interface-name</i> gigheter-options 802.3ad aex</b></pre>      |
| <p>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.</p>             | <pre>[edit interfaces] user@host# <b>set aex aggregated-ether-options link- speed <i>speed</i></b></pre>    |
| <p>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>.</p> <p>You cannot configure the minimum number of links and the minimum bandwidth at the same time. They are mutually exclusive.</p> | <pre>[edit interfaces] user@host# <b>set aex aggregated-ether-options minimum-links <i>number</i></b></pre> |
| <p>Step 5: (Optional) Specify the minimum bandwidth for the aggregated Ethernet links.</p> <p>You cannot configure link protection with minimum bandwidth.</p> <p>You cannot configure the minimum number of links and the minimum bandwidth at the same time. They are mutually exclusive.</p>                                                                     | <pre>[edit interfaces] user@host# <b>set aex aggregated-ether-options minimum-bandwidth</b></pre>           |

Table 8: Aggregated Ethernet Interfaces Configuration (Continued)

| Configuration Step                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Command                                                                                                                                                                                                                            |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 6: Specify an interface family and the IP address for the aggregated Ethernet bundle. Aggregated Ethernet interfaces can be VLAN-tagged or untagged.</p> <p>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.</p> | <p>Tagged Interface</p> <pre>[edit interfaces] user@host# set aex vlan-tagging unit 0 vlan-id <i>vlan-id</i></pre> <p>Untagged Interface</p> <pre>[edit interfaces] user@host# set aex unit 0 family inet address ip-address</pre> |
| <p>Step 7: (Optional) Configure your device to collect multicast statistics for the aggregated Ethernet interface.</p> <p>To view the multicast statistics, use the <b>show interfaces statistics detail</b> command. If you have not configured collection of multicast statistics, you cannot view the multicast statistics.</p>                                                                                                                                                     | <pre>[edit interfaces] user@host# set aex multicast-statistics</pre>                                                                                                                                                               |
| <p>Step 8: Verify and commit the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                    | <pre>[edit interfaces] user@host# run show configuration user@host# commit</pre>                                                                                                                                                   |
| <p>Step 9: (Optional) Delete an aggregated Ethernet Interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                     | <pre>[edit] user@host# delete interfaces aex</pre> <p>OR</p> <pre>[edit] user@host# delete chassis aggregated-devices ethernet device-count</pre>                                                                                  |

## SEE ALSO

[802.3ad](#)

[device-count](#) | [837](#)

[link-speed \(Aggregated Ethernet\)](#) | [969](#)

[minimum-bandwidth \(aggregated Ethernet\)](#) | [1010](#)

[minimum-links](#) | [1012](#)

[multicast-statistics](#) | [1021](#)

## Mixed-Mode and Mixed-Rate Aggregated Ethernet Interfaces

### IN THIS SECTION

- [Benefits](#) | [67](#)

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 completely 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 9 on page 68](#) lists the platforms and corresponding MPCs that support mixed-rate aggregated Ethernet bundles on MX Series routers.

**Table 9: 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          |
| 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-MPC2E-3D-Q; MX-MPC2-3D-EQ; MX-MPC2E-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          |

**Table 9: Platform Support Matrix for Mixed-Rate Aggregated Ethernet Bundles on MX Series Routers**  
(Continued)

| Supported MPCs            | Supported Platform      | Initial Release |
|---------------------------|-------------------------|-----------------|
| MPC10E (MPC10E-15C-MRATE) | MX240, MX480, and MX960 | 19.1R1          |

Table 10 on page 69 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

**Table 10: 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<br><br>(WAN rate: OC192) | T640, T1600, T4000, and TX Matrix Plus routers | <ul style="list-style-type: none"> <li>T4000 FPC5 (T4000-FPC5-3D)</li> </ul>                                                                                                                    | <ul style="list-style-type: none"> <li>10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (PF-24XGE-SFPP)</li> <li>10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)</li> </ul> |
|                                                          |                                                | <ul style="list-style-type: none"> <li>Enhanced Scaling FPC3 (T640-FPC3-ES)</li> </ul>                                                                                                          | <ul style="list-style-type: none"> <li>10-Gigabit Ethernet PIC with XENPAK (PC-1XGE-XENPAK)</li> </ul>                                                                                              |
|                                                          |                                                | <ul style="list-style-type: none"> <li>Enhanced Scaling FPC4 (T640-FPC4-ES)</li> <li>Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES)</li> <li>T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)</li> </ul> | <ul style="list-style-type: none"> <li>10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP)</li> <li>10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)</li> </ul>                       |

**Table 10: Platform Support Matrix for Mixed Aggregated Ethernet Bundles on T Series (Continued)**

| Rate and Mode                             | Supported Platform                             | Supported FPCs                                                                                                                                                                                  | Supported PICs                                                                                                                                                                                                                                                                                            |
|-------------------------------------------|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 40-Gigabit Ethernet, 100-Gigabit Ethernet | T4000 and TX Matrix Plus routers               | <ul style="list-style-type: none"> <li>T4000 FPC5 (T4000-FPC5-3D)</li> </ul>                                                                                                                    | <ul style="list-style-type: none"> <li>100-Gigabit Ethernet PIC with CFP (PF-1CGE-CFP)</li> </ul>                                                                                                                                                                                                         |
|                                           | T640, T1600, T4000, and TX Matrix Plus routers | <ul style="list-style-type: none"> <li>Enhanced Scaling FPC4 (T640-FPC4-ES)</li> <li>Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES)</li> <li>T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES)</li> </ul> | <ul style="list-style-type: none"> <li>100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4)</li> </ul> <p><b>NOTE:</b> This PIC is available packaged only in an assembly with the T1600-FPC4-ES FPC.</p> <ul style="list-style-type: none"> <li>40-Gigabit Ethernet PIC with CFP (PD-1XLE-CFP)</li> </ul> |

### 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 11 on page 71](#) describes the steps to configure mixed-rate aggregated Ethernet bundle on your device.

**Table 11: 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 <b>device-count</b> value as 2, you can configure two aggregated bundles. | <pre>[edit chassis aggregated-devices ethernet] user@host# <b>set device-count <i>number</i></b></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# <b>set <i>interface-name</i> together-options 802.3ad ae<i>x</i></b></pre> |

Table 11: Mixed-Rate Aggregated Ethernet Configuration (*Continued*)

| Configuration Step                                                                                                                                                                                                                                                                                                                                                                                                                      | Command                                                                                    |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| <p>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.</p> <p>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.</p> | <pre>[edit interfaces] user@host# set aex aggregated-ether-options link-speed mixed</pre>  |
| <p>Step 4: Specify the minimum bandwidth for the aggregated Ethernet links.</p> <p>You cannot configure link protection with the minimum bandwidth.</p>                                                                                                                                                                                                                                                                                 | <pre>[edit interfaces] user@host# set aex aggregated-ether-options minimum-bandwidth</pre> |
| <p>Step 5: Verify and commit the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                     | <pre>[edit interfaces] user@host# run show configuration user@host# commit</pre>           |

## SEE ALSO

[802.3ad](#)

[device-count | 837](#)

[link-speed \(Aggregated Ethernet\) | 969](#)

[minimum-bandwidth \(aggregated Ethernet\) | 1010](#)

## What Is Link Aggregation Control Protocol?

### IN THIS SECTION

- [Benefits | 73](#)

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 (00:00:00:00:00:00), 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.

## Configure LACP

Table 12 on page 74 describes the steps to configure LACP on an aggregated Ethernet interface.

**Table 12: LACP Configuration**

| Configuration Step                                                                                                                                                                                                                                                                                                                                                                                                                               | Command                                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 1: Specify the LACP transmission mode - active or passive.</p>                                                                                                                                                                                                                                                                                                                                                                           | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options] user@host# <b>set lacp active</b> user@host# <b>set lacp passive</b></pre> |
| <p>Step 2: Specify the interval at which the interfaces send LACP packets.</p> <p>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.</p>                                                                                                                                                                             | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options lacp] user@host# <b>set periodic <i>interval</i></b></pre>                  |
| <p>Step 3: Configure the LACP system identifier.</p> <p>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.</p> <p>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.</p>                                       | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options lacp] user@host# <b>set system-id <i>system-id</i></b></pre>                |
| <p>Step 4: Configure the LACP system priority at the Aggregated Ethernet interface level.</p> <p>This system priority takes precedence over the priority value configured at the global <b>[edit chassis]</b> 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.</p> | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options lacp] user@host# <b>set system-priority <i>system-priority</i></b></pre>    |

Table 12: LACP Configuration (*Continued*)

| Configuration Step                                                                                                                                                                                                                                                                 | Command                                                                                                                                            |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 5: (Optional) Configure the LACP administrative key.</p> <p>You must configure MC-LAG to configure this option. For more information on MC-LAG, see <a href="#">Understanding Multichassis Link Aggregation Groups</a>.</p>                                                | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options lacp] user@host# <b>set admin-key <i>number</i></b></pre>                     |
| <p>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.</p> | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options lacp] user@host# <b>set hold-time <i>timer-value</i></b></pre>                |
| <p>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.</p>                                                                                                                          | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options lacp] user@host# <b>set accept-data</b></pre>                                 |
| <p>Step 8: Verify and commit the configuration.</p>                                                                                                                                                                                                                                | <pre>[edit interfaces <i>interface-name</i> aggregated-ether-options lacp] user@host# <b>run show configuration</b> user@host# <b>commit</b></pre> |

## SEE ALSO

---

[hold-time up | 907](#)


---

[periodic | 1069](#)


---

[system-id | 1164](#)


---

[system-priority | 1167](#)

## Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links

### IN THIS SECTION

- [Benefits | 77](#)

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 | 81
- Overview | 81
- Verification | 82

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 gigheter-options 802.3ad ae10
distribution-list dl1
user@host# set interfaces ge-0/0/4 apply-groups-except INTF gigheter-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 gigheter-options 802.3ad ae10
distribution-list dl1
[edit]
user@host# set interfaces ge-0/0/4 apply-groups-except INTF gigheter-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 *d/1* as the primary list, and those ending in an even number are assigned the distribution list *d/2* 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 | 82](#)

### *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 : 0
 Input packets : 0
 Output packets : 0
 Dropped traffic statistics due to STP State:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets : 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed
discards: 0, Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
 Ingress queues: 8 supported, 4 in use
 Queue counters: Queued packets Transmitted packets Dropped packets
 0 0 0 0
```

```

1 0 0 0
2 0 0 0
3 0 0 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
0 0 0 0
1 0 0 0
2 0 0 0
3 0 0 0
Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control

```

Logical interface ae10.101 (Index 345) (SNMP ifIndex 617) (Generation 154)

Description: matched odd

Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.101 ] Encapsulation: ENET2

| Statistics | Packets | pps | Bytes | bps |
|------------|---------|-----|-------|-----|
|------------|---------|-----|-------|-----|

Bundle:

|         |   |   |    |   |
|---------|---|---|----|---|
| Input : | 0 | 0 | 0  | 0 |
| Output: | 2 | 0 | 92 | 0 |

Adaptive Statistics:

|                   |   |
|-------------------|---|
| Adaptive Adjusts: | 0 |
| Adaptive Scans :  | 0 |
| Adaptive Updates: | 0 |

Link:

ge-0/0/3.101

|         |   |   |    |   |
|---------|---|---|----|---|
| Input : | 0 | 0 | 0  | 0 |
| Output: | 2 | 0 | 92 | 0 |

ge-0/0/4.101

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
| Output: | 0 | 0 | 0 | 0 |

Aggregate member links: 2

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

| List-Type | Status |
|-----------|--------|
|-----------|--------|

|         |        |
|---------|--------|
| Primary | Active |
|---------|--------|

```

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

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

Logical interface ae10.102 (Index 344) (SNMP ifIndex 615) (Generation 153)
Description: matched even
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.102] Encapsulation: ENET2
Statistics Packets pps Bytes bps
Bundle:
 Input : 0 0 0 0
 Output: 4 0 296 0
Adaptive Statistics:
 Adaptive Adjusts: 0
 Adaptive Scans : 0
 Adaptive Updates: 0
Link:
 ge-0/0/3.102
 Input : 0 0 0 0
 Output: 4 0 296 0
 ge-0/0/4.102
 Input : 0 0 0 0
 Output: 0 0 0 0
Marker Statistics: Marker Rx Resp Tx Unknown Rx Illegal Rx
 ge-0/0/3.102 0 0 0 0
 ge-0/0/4.102 0 0 0 0

List-Type Status
Primary Active
 Interfaces:
 ge-0/0/3 Up

```

```

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

Protocol inet, MTU: 8978, Generation: 196, Route table: 0
 Flags: Sendbcst-pkt-to-re
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.2.0.1 , Local: 10.2.0.1, Broadcast: 10.2.0.3,
Generation: 152
Protocol multiservice, MTU: Unlimited, Generation: 197, Route table: 0
 Policer: Input: __default_arp_policer__

Logical interface ae10.103 (Index 343) (SNMP ifIndex 614) (Generation 152)
Description: matched odd
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.103] Encapsulation: ENET2
Statistics Packets pps Bytes bps
Bundle:
 Input : 0 0 0 0
 Output: 3 0 194 0
Adaptive Statistics:
 Adaptive Adjusts: 0
 Adaptive Scans : 0
 Adaptive Updates: 0
Link:
 ge-0/0/3.103
 Input : 0 0 0 0
 Output: 3 0 194 0
 ge-0/0/4.103
 Input : 0 0 0 0
 Output: 0 0 0 0
Marker Statistics: Marker Rx Resp Tx Unknown Rx Illegal Rx
 ge-0/0/3.103 0 0 0 0
 ge-0/0/4.103 0 0 0 0

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

```

```

 Interfaces:
 ge-0/0/3 Up
List-Type Status
Standby Down

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

Logical interface ae10.104 (Index 342) (SNMP ifIndex 616) (Generation 151)
Description: matched even
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.104] Encapsulation: ENET2
Statistics Packets pps Bytes bps
Bundle:
 Input : 0 0 0 0
 Output: 2 0 92 0
Adaptive Statistics:
 Adaptive Adjusts: 0
 Adaptive Scans : 0
 Adaptive Updates: 0
Link:
 ge-0/0/3.104
 Input : 0 0 0 0
 Output: 2 0 92 0
 ge-0/0/4.104
 Input : 0 0 0 0
 Output: 0 0 0 0
Marker Statistics: Marker Rx Resp Tx Unknown Rx Illegal Rx
 ge-0/0/3.104 0 0 0 0
 ge-0/0/4.104 0 0 0 0

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

```

```

List-Type Status
Standby Down

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

Logical interface ae10.32767 (Index 341) (SNMP ifIndex 613) (Generation 150)
 Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [0x0000.0] Encapsulation: ENET2
 Statistics Packets pps Bytes bps
Bundle:
 Input : 0 0 0 0
 Output: 0 0 0 0
Adaptive Statistics:
 Adaptive Adjusts: 0
 Adaptive Scans : 0
 Adaptive Updates: 0
Link:
 ge-0/0/3.32767
 Input : 0 0 0 0
 Output: 95 0 38039 0
 ge-0/0/4.32767
 Input : 0 0 0 0
 Output: 95 0 38039 0
Marker Statistics: Marker Rx Resp Tx Unknown Rx Illegal Rx
 ge-0/0/3.32767 0 0 0 0
 ge-0/0/4.32767 0 0 0 0
Protocol multiservice, MTU: Unlimited, Generation: 191, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__

```

## SEE ALSO

[distribution-list](#)

[targeted-distribution](#) | 1173

## Independent Micro-BFD Sessions for LAG

### IN THIS SECTION

- [Benefits](#) | 90

The Bidirectional Forwarding Detection (BFD) protocol is a simple detection protocol that quickly detects failures in the forwarding paths. To enable failure detection for aggregated Ethernet interfaces in a LAG, you can configure an independent, asynchronous-mode BFD session on every LAG member link in a LAG bundle. Instead of a single BFD session monitoring the status of the UDP port, independent micro-BFD sessions monitor the status of individual member links.

When you configure micro-BFD sessions on every member link in a LAG bundle, each individual session determines the Layer 2 and Layer 3 connectivity of each member link in a LAG.

After the individual session is established on a particular link, member links are attached to the LAG and then load balanced by either one of the following:

- Static configuration—The device control process acts as the client to the micro-BFD session.
- Link Aggregation Control Protocol (LACP)—LACP acts as the client to the micro-BFD session.

When the micro-BFD session is up, a LAG link is established and data is transmitted over that LAG link. If the micro-BFD session on a member link is down, that particular member link is removed from the load balancer, and the LAG managers stop directing traffic to that link. These micro-BFD sessions are independent of each other despite having a single client that manages the LAG interface.

Micro-BFD sessions run in the following modes:

- Distribution mode—In this mode, the Packet Forwarding Engine (PFE) sends and receives the packets at Layer 3. By default, micro-BFD sessions are distributed at Layer 3.
- Non-distribution mode—In this mode, the Routing Engine sends and receives the packets at Layer 2. You can configure the BFD session to run in this mode by including the **no-delegate-processing** statement under periodic packet management (PPM).

A pair of routing devices in a LAG exchange BFD packets at a specified, regular interval. The routing device detects a neighbor failure when it stops receiving a reply after a specified interval. This allows the quick verification of member link connectivity with or without LACP. A UDP port distinguishes BFD over LAG packets from BFD over single-hop IP packets. The Internet Assigned Numbers Authority (IANA) has allocated 6784 as the UDP destination port for micro-BFD.

## Benefits

- Failure detection for LAG—Enables failure detection between devices that are in point-to-point connections.
- Multiple BFD sessions—Enables you to configure multiple micro-BFD sessions for each member link instead of a single BFD session for the entire bundle.

## Configuration Guidelines for Micro-BFD Sessions

Consider the following guidelines as you configure individual micro-BFD sessions on an aggregated Ethernet bundle.

- This feature works only when both the devices support BFD. If BFD is configured at one end of the LAG, this feature does not work.
- Starting with Junos OS Release 13.3, IANA has allocated 01-00-5E-90-00-01 as the dedicated MAC address for micro BFD. Dedicated MAC mode is used by default for micro BFD sessions.
- In Junos OS, micro-BFD control packets are always untagged by default. For Layer 2 aggregated interfaces, the configuration must include **vlan-tagging** or **flexible-vlan-tagging** options when you configure Aggregated Ethernet with BFD. Otherwise, the system will throw an error while committing the configuration.
- When you enable micro-BFD on an aggregated Ethernet interface, the aggregated interface can receive micro-BFD packets. 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.
- 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. 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.
- 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. Junos OS performs this check on both IPv4 and IPv6 micro-BFD address configurations, and if they do not match, the commit fails.
- For the IPv6 address family, disable duplicate address detection before configuring this feature with aggregated Ethernet interface addresses. To disable duplicate address detection, include the **dad-disable** statement at the `[edit interface aexunit y family inet6]` hierarchy level.



**CAUTION:** Deactivate `bfd-liveness-detection` at the **[edit interfaces aex aggregated-ether-options]** hierarchy level or deactivate the aggregated Ethernet interface before changing the neighbor address from the loopback IP address to the aggregated Ethernet interface IP address. Modifying the local and neighbor address without deactivating **bfd-liveness-detection** or the aggregated Ethernet interface first might cause micro-BFD sessions failure.

## Example: Configure Independent Micro-BFD Sessions for LAG

### IN THIS SECTION

- [Requirements | 91](#)
- [Overview | 92](#)
- [Configuration | 92](#)
- [Verification | 100](#)

This example shows how to configure an independent micro-BFD session for aggregated Ethernet interfaces.

### Requirements

This example uses the following hardware and software components:

- MX Series routers with Junos MPCs
- T Series routers with Type 4 FPC or Type 5 FPC

BFD for LAG is supported on the following PIC types on T-Series:

- PC-1XGE-XENPAK (Type 3 FPC)
- PD-4XGE-XFP (Type 4 FPC)
- PD-5-10XGE-SFPP (Type 4 FPC)
- 24x10GE (LAN/WAN) SFPP, 12x10GE (LAN/WAN) SFPP, 1X100GE Type 5 PICs
- PTX Series routers with 24X10GE (LAN/WAN) SFPP
- Junos OS Release 13.3 or later running on all devices

## Overview

### IN THIS SECTION


- [Topology | 92](#)

The example includes two routers that are directly connected. Configure two aggregated Ethernet interfaces, AE0 for IPv4 connectivity and AE1 for IPv6 connectivity. Configure a micro-BFD session on the AE0 bundle using IPv4 addresses as local and neighbor endpoints on both routers. Configure a micro-BFD session on the AE1 bundle using IPv6 addresses as local and neighbor endpoints on both routers. This example verifies that independent micro-BFD sessions are active in the output.

### Topology

[Figure 2 on page 92](#) shows the sample topology.

**Figure 2: Configuring an Independent Micro-BFD Session for LAG**



## Configuration

### IN THIS SECTION

- [CLI Quick Configuration | 93](#)
- [Configure a Micro-BFD Session for Aggregated Ethernet Interfaces | 94](#)
- [Procedure | 94](#)
- [Results | 97](#)

### 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.

#### Router R0

```
set interfaces ge-1/0/1 unit 0 family inet address 20.20.20.1/30
set interfaces ge-1/0/1 unit 0 family inet6 address 3ffe::1:1/126
set interfaces xe-4/0/0 gigether-options 802.3ad ae0
set interfaces xe-4/0/1 gigether-options 802.3ad ae0
set interfaces xe-4/1/0 gigether-options 802.3ad ae1
set interfaces xe-4/1/1 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.107/32
set interfaces lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/126
set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval 100
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor 10.255.106.102
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address 10.255.106.107
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 10.0.0.1/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval 100
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor 201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address 201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet6 address 5555::1/126
set interface ae1 unit 0 family inet6 dad-disable
set routing-options nonstop-routing
set routing-options static route 30.30.30.0/30 next-hop 10.0.0.2
set routing-options rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2
set protocols bfd traceoptions file bfd
set protocols bfd traceoptions file size 100m
set protocols bfd traceoptions file files 10
set protocols bfd traceoptions flag all
```

## Router R1

```

set interfaces ge-1/1/8 unit 0 family inet address 30.30.30.1/30
set interfaces ge-1/1/8 unit 0 family inet6 address 3ffe::1:2/126
set interfaces xe-0/0/0 gigether-options 802.3ad ae0
set interfaces xe-0/0/1 gigether-options 802.3ad ae0
set interfaces xe-0/0/2 gigether-options 802.3ad ae1
set interfaces xe-0/0/3 gigether-options 802.3ad ae1
set interfaces lo0 unit 0 family inet address 10.255.106.102/32
set interfaces lo0 unit 0 family inet6 address 201:DB8:251::bb:bb:1/126
set interfaces ae0 aggregated-ether-options bfd-liveness-detection minimum-interval 150
set interfaces ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae0 aggregated-ether-options bfd-liveness-detection neighbor 10.255.106.107
set interfaces ae0 aggregated-ether-options bfd-liveness-detection local-address 10.255.106.102
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 family inet address 10.0.0.2/30
set interfaces ae1 aggregated-ether-options bfd-liveness-detection minimum-interval 200
set interfaces ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
set interfaces ae1 aggregated-ether-options bfd-liveness-detection neighbor 201:DB8:251::aa:aa:1
set interfaces ae1 aggregated-ether-options bfd-liveness-detection local-address 201:DB8:251::bb:bb:1
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options lacp passive
set interfaces ae1 unit 0 family inet6 address 5555::2/126
set routing-options static route 20.20.20.0/30 next-hop 10.0.0.1
set routing-options rib inet6.0 static route 3ffe::1:1/126 next-hop 5555::1

```

### *Configure a Micro-BFD Session for Aggregated Ethernet Interfaces*

#### *Procedure*

#### **Step-by-Step Procedure**

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see “*Using the CLI Editor in Configuration Mode*” in the [CLI User Guide](#).

**NOTE:** Repeat this procedure for Router R1, modifying the appropriate interface names, addresses, and any other parameters for each router.

To configure a micro-BFD session for aggregated Ethernet interfaces on Router R0:

**1.** Configure the physical interfaces.

```
[edit interfaces]
user@R0# set ge-1/0/1 unit 0 family inet address 20.20.20.1/30
user@R0# set ge-1/0/1 unit 0 family inet6 address 3ffe::1:1/126
user@R0# set xe-4/0/0 gigether-options 802.3ad ae0
user@R0# set xe-4/0/1 gigether-options 802.3ad ae0
user@R0# set xe-4/1/0 gigether-options 802.3ad ae1
user@R0# set xe-4/1/1 gigether-options 802.3ad ae1
```

**2.** Configure the loopback interface.

```
[edit interfaces]
user@R0# set lo0 unit 0 family inet address 10.255.106.107/32
user@R0# set lo0 unit 0 family inet6 address 201:DB8:251::aa:aa:1/128
```

**3.** Configure an IP address on the aggregated Ethernet interface ae0 with either IPv4 or IPv6 addresses, according to your network requirements.

```
[edit interfaces]
user@R0# set ae0 unit 0 family inet address 10.0.0.1/30
```

**4.** Set the routing option, create a static route, and set the next-hop address.

**NOTE:** You can configure either an IPv4 or IPv6 static route, depending on your network requirements.

```
[edit routing-options]
user@R0# set nonstop-routing
```

```

user@R0# set static route 30.30.30.0/30 next-hop 10.0.0.2
user@R0# set rib inet6.0 static route 3ffe::1:2/126 next-hop 5555::2

```

5. Configure Link Aggregation Control Protocol (LACP).

```

[edit interfaces]
user@R0# set ae0 aggregated-ether-options lacp active

```

6. Configure BFD for the aggregated Ethernet interface ae0, and specify the minimum interval, local IP address, and the neighbor IP address.

```

[edit interfaces]
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection minimum-interval 100
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection multiplier 3
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection neighbor 10.255.106.102
user@R0# set ae0 aggregated-ether-options bfd-liveness-detection local-address 10.255.106.107
user@R0# set ae0 aggregated-ether-options minimum-links 1
user@R0# set ae0 aggregated-ether-options link-speed 10g

```

7. Configure an IP address on the aggregated Ethernet interface ae1.

You can assign either IPv4 or IPv6 addresses as per your network requirements.

```

[edit interfaces]
user@R0# set ae1 unit 0 family inet6 address 5555::1/126

```

8. Configure BFD for the aggregated Ethernet interface ae1.

```

[edit interfaces]
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection minimum-interval 100
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection multiplier 3
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection neighbor 201:DB8:251::bb:bb:1
user@R0# set ae1 aggregated-ether-options bfd-liveness-detection local-address 201:DB8:251::aa:aa:1
user@R0# set ae1 aggregated-ether-options minimum-links 1
user@R0# set ae1 aggregated-ether-options link-speed 10g

```

**NOTE:** Starting with Junos OS Release 16.1, you can also configure this feature with the aggregated Ethernet interface address as the local address in a micro-BFD session.

Starting 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. Junos OS performs this check on both IPv4 and IPv6 micro-BFD address configurations, and if they do not match, the commit fails.

## 9. Configure tracing options for BFD for troubleshooting.

```
[edit protocols]
user@R0# set bfd traceoptions file bfd
user@R0# set bfd traceoptions file size 100m
user@R0# set bfd traceoptions file files 10
user@R0# set bfd traceoptions flag all
```

### Results

From operational mode, enter the **show interfaces**, **show protocols**, and **show routing-options** commands and confirm your configuration. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0> show interfaces
traceoptions {
 flag bfd-events;
}
ge-1/0/1 {
 unit 0 {
 family inet {
 address 20.20.20.1/30;
 }
 family inet6 {
 address 3ffe::1:1/126;
 }
 }
}
xe-4/0/0 {
 enable;
 gigether-options {
```

```
 802.3ad ae0;
 }
}
xe-4/0/1 {
 gigether-options {
 802.3ad ae0;
 }
}
xe-4/1/0 {
 enable;
 gigether-options {
 802.3ad ae1;
 }
}
xe-4/1/1 {
 gigether-options {
 802.3ad ae1;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 10.255.106.107/32;
 }
 family inet6 {
 address 201:DB8:251::aa:aa:1/128;
 }
 }
}
ae0 {
 aggregated-ether-options {
 bfd-liveness-detection {
 minimum-interval 100;
 neighbor 10.255.106.102;
 local-address 10.255.106.107;
 }
 minimum-links 1;
 link-speed 10g;
 lacp {
 active;
 }
 }
 unit 0 {
```

```
 family inet {
 address 10.0.0.1/30;
 }
 }
}
ae1 {
 aggregated-ether-options {
 bfd-liveness-detection {
 minimum-interval 100;
 multiplier 3;
 neighbor 201:DB8:251::bb:bb:1;
 local-address 201:DB8:251::aa:aa:1;
 }
 minimum-links 1
 link-speed 10g;
 }
 unit 0 {
 family inet6 {
 address 5555::1/126;
 }
 }
}
}
```

```
user@R0> show protocols
bfd {
 traceoptions {
 file bfd size 100m files 10;
 flag all;
 }
}
}
```

```
user@R0> show routing-options
nonstop-routing ;
rib inet6.0 {
 static {
 route 3ffe:1:2/126 {
 next-hop 5555::2;
 }
 }
}
}
```

```
static {
 route 30.30.30.0/30 {
 next-hop 10.0.0.2;
 }
}
```

If you are done configuring the device, commit the configuration.

```
user@R0# commit
```

## Verification

### IN THIS SECTION

- [Verify That the Independent BFD Sessions Are Up | 100](#)
- [View Detailed BFD Events | 102](#)

Confirm that the configuration is working properly.

### *Verify That the Independent BFD Sessions Are Up*

## Purpose

Verify that the micro-BFD sessions are up, and view details about the BFD sessions.

## Action

From operational mode, enter the **show bfd session extensive** command.

```
user@R0> show bfd session extensive
```

| Address        | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------------|-------|-----------|-------------|-------------------|------------|
| 10.255.106.102 | Up    | xe-4/0/0  | 9.000       | 3.000             |            |

```
3
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13, previous down time 00:00:06
```

Local diagnostic None, remote diagnostic None  
 Remote heard, hears us, version 1  
 Replicated  
 Session type: **Micro BFD**  
 Min async interval 0.100, min slow interval 1.000  
 Adaptive async TX interval 0.100, RX interval 0.100  
 Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3  
 Remote min TX interval 3.000, min RX interval 3.000, multiplier 3  
 Local discriminator 21, remote discriminator 75  
 Echo mode disabled/inactive  
 Remote is control-plane independent  
 Session ID: 0x0

| Address        | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------------|-------|-----------|-------------|-------------------|------------|
| 10.255.106.102 | Up    | xe-4/0/1  | 9.000       | 3.000             |            |

3

Client LACPD, TX interval 0.100, RX interval 0.100  
 Session up time 4d 23:13, previous down time 00:00:07  
 Local diagnostic None, remote diagnostic None  
 Remote heard, hears us, version 1  
 Replicated  
 Session type: **Micro BFD**  
 Min async interval 0.100, min slow interval 1.000  
 Adaptive async TX interval 0.100, RX interval 0.100  
 Local min TX interval 0.100, minimum RX interval 0.100, multiplier 3  
 Remote min TX interval 3.000, min RX interval 3.000, multiplier 3  
 Local discriminator 19, remote discriminator 74  
 Echo mode disabled/inactive  
 Remote is control-plane independent  
 Session ID: 0x0

| Address              | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------------------|-------|-----------|-------------|-------------------|------------|
| 201:DB8:251::bb:bb:1 | Up    | xe-4/1/1  | 9.000       |                   |            |

3.000 3

Client LACPD, TX interval 0.100, RX interval 0.100  
 Session up time 4d 23:13  
 Local diagnostic None, remote diagnostic None  
 Remote not heard, hears us, version 1  
 Replicated  
 Session type: **Micro BFD**  
 Min async interval 0.100, min slow interval 1.000

```

Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 17, remote discriminator 67
Echo mode disabled/inactive, no-absorb, no-refresh
Remote is control-plane independent
Session ID: 0x0

Address State Interface Detect Transmit
 Time Interval Multiplier
201:DB8:251::bb:bb:1 UP xe-4/1/0 9.000
3.000 3
Client LACPD, TX interval 0.100, RX interval 0.100
Session up time 4d 23:13
Local diagnostic None, remote diagnostic None
Remote not heard, hears us, version 1
Replicated
Session type: Micro BFD
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 3.000, min RX interval 3.000, multiplier 3
Local discriminator 16, remote discriminator 66
Echo mode disabled/inactive, no-absorb, no-refresh
Remote is control-plane independent
Session ID: 0x0

4 sessions, 4 clients
Cumulative transmit rate 2.0 pps, cumulative receive rate 1.7 pps

```

## Meaning

The Session Type field represents the independent micro-BFD sessions running on the links in a LAG. The TX interval *value*, RX interval *value* output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

### *View Detailed BFD Events*

## Purpose

View the contents of the BFD trace file to assist in troubleshooting, if required.

## Action

From operational mode, enter the `file show /var/log/bfd` command.

```

user@R0> file show /var/log/bfd
Jun 5 00:48:59 Protocol (1) len 1: BFD
Jun 5 00:48:59 Data (9) len 41: (hex) 42 46 44 20 6e 65 69 67 68 62 6f 72 20
31 30 2e 30 2e 30
Jun 5 00:48:59 PPM Trace: BFD neighbor 10.255.106.102 (IFL 349) set, 9 0
Jun 5 00:48:59 Received Downstream RcvPkt (19) len 108:
Jun 5 00:48:59 IfIndex (3) len 4: 329
Jun 5 00:48:59 Protocol (1) len 1: BFD
Jun 5 00:48:59 SrcAddr (5) len 8: 10.255.106.102
Jun 5 00:48:59 Data (9) len 24: (hex) 00 88 03 18 00 00 00 4b 00 00 00 15 00
2d c6 c0 00 2d c6
Jun 5 00:48:59 PktError (26) len 4: 0
Jun 5 00:48:59 RtblIdx (24) len 4: 0
Jun 5 00:48:59 MultiHop (64) len 1: (hex) 00
Jun 5 00:48:59 Unknown (168) len 1: (hex) 01
Jun 5 00:48:59 Unknown (171) len 2: (hex) 02 3d
Jun 5 00:48:59 Unknown (172) len 6: (hex) 80 71 1f c7 81 c0
Jun 5 00:48:59 Authenticated (121) len 1: (hex) 01
Jun 5 00:48:59 BFD packet from 10.0.0.2 (IFL 329), len 24
Jun 5 00:48:59 Ver 0, diag 0, mult 3, len 24
Jun 5 00:48:59 Flags: IHU Fate
Jun 5 00:48:59 My discr 0x0000004b, your discr 0x00000015
Jun 5 00:48:59 Tx ivl 3000000, rx ivl 3000000, echo rx ivl 0
Jun 5 00:48:59 [THROTTLE]bfd_rate_limit_can_accept_pkt: session 10.255.106.102
is up or already in program thread
Jun 5 00:48:59 Replicate: marked session (discr 21) for update

```

## Meaning

BFD messages are being written to the specified trace file.

## SEE ALSO

[authentication](#)

*bfd-liveness-detection*

## MAC Address Accounting for Dynamically Learned Addresses on Aggregated Ethernet Interfaces

### IN THIS SECTION

- [Benefits](#) | 105

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 | 105](#)

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.

### Release History Table

| 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 | 106](#)
- [Configuring Aggregated Ethernet Minimum Links | 109](#)
- [Example: Configuring Aggregated Ethernet Link Protection | 110](#)

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

- [Configuring Link Protection for Aggregated Ethernet Interfaces | 107](#)
- [Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces | 108](#)
- [Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link | 108](#)

- [Disabling Link Protection for Aggregated Ethernet Interfaces](#) | 109

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](#) | 964

*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 | ggether-options) 802.3ad aex primary
```

2. Configure the backup logical interface.

```
[edit interfaces interface-name]
user@host# set (fastether-options | ggether-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 interface revert aex** 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](#) | 155

# Scheduling on Aggregated Ethernet Interfaces

## IN THIS SECTION

- [Configuring Shared Scheduling on Aggregated Ethernet Interfaces](#) | 111
- [Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection](#) | 112

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 aeX]** 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 aeX]** hierarchy level. The default setting is **scale**.

```
[edit class-of-service interfaces aeX 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 | 163](#)

[Performance Monitoring on Aggregated Ethernet Interfaces | 155](#)

## Load Balancing on Aggregated Ethernet Interfaces

### IN THIS SECTION

- [Load Balancing and Ethernet Link Aggregation Overview | 114](#)
- [Understanding Aggregated Ethernet Load Balancing | 115](#)
- [Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data | 117](#)

- [Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | 120](#)
- [Configuring Adaptive Load Balancing | 121](#)
- [Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers | 123](#)
- [Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers | 130](#)
- [Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers | 133](#)
- [Example: Configuring Aggregated Ethernet Load Balancing | 136](#)

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.)

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input *logical interface* (unit). The default LAG algorithm is optimized for Layer 2 switching. 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. 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.

### SEE ALSO

| [payload](#)

## Understanding Aggregated Ethernet Load Balancing

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

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

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

Typically, the applications that use load balancing include:

- Aggregated Interfaces (Layer 2)

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

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

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

ECMP load balancing can be:

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

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

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

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

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

- Adaptive Load Balancing

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

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

**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 Aggregated Ethernet Interfaces or LAG Bundles | 119](#)

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 **ae**bundle.

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

$$V_i = ((C_i - (M/N))/N)$$

where

- $V_i$  denotes the deviation for that child link 'i'.
- $i$  denotes the child link member/index.
- $C_i$  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 unicast 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 aeX unit logical-unit-number forwarding-options load-balance-stateful]** hierarchy level. By default, Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. All Packet Forwarding Engine slots are assigned the same hash value by default. To configure the load-balancing algorithm to dynamically rebalance the

LAG using existing parameters, include the **rebalance *interval*** statement at the **[edit interfaces aeX unit *logical-unit-number forwarding-options load-balance-stateful*]** hierarchy level. This parameter periodically load balances traffic by providing a synchronized rebalance switchover across all the ingress Packet Forwarding Engines (PFEs) over a rebalance interval. You can specify the interval as a value in the range of 1 through 1000 flows per minute. To configure the load type, include the **load-type (low | medium | high)** statement at the **[edit interfaces aeX unit *logical-unit-number forwarding-options load-balance-stateful*]** hierarchy level.

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

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

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

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

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

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

- This capability is not supported for multicast traffic (native/flood).
- Enabling the rebalance option or clearing the load balance state can cause packet reordering for active flows because different sets of links can be selected for traffic flows.

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

## SEE ALSO

[Link Protection of Aggregated Ethernet Interfaces](#) | 106

## 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 aeXunit logical-unit-number
```

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

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

- 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 aeXunit logical-unit-number load-balance-stateful]
user@R2# set per-flow
```

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

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

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

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

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

## SEE ALSO

[Link Protection of Aggregated Ethernet Interfaces](#) | 106

## 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 | 123](#)
- [Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers | 123](#)
- [Configuring Symmetrical Load Balancing on Trio-Based MPCs | 127](#)
- [Example Configurations | 129](#)

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

MX Series routers with Aggregated Ethernet PICs support symmetrical load balancing on an 802.3ad LAG. This feature is significant when two MX Series routers are connected transparently through deep packet inspection (DPI) devices over an LAG bundle. DPI devices keep track of flows and require information of a given flow in both forward and reverse directions. Without symmetrical load balancing on an 802.3ad LAG, the DPIs could misunderstand the flow, leading to traffic disruptions. By using this feature, a given flow of traffic (duplex) is ensured for the same devices in both directions.

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


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

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

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

Consider the example in [Figure 3 on page 124](#).

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



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

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

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

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

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

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

### Example Configuration Statements

To configure 802.3ad LAG parameters at the bundle level:

```
[edit interfaces]
g(x)e-fpc/pic/port {
 together-options {
 802.3ad {
 bundle;
 link-index number;
 }
 }
}
```

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

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

```
[edit forwarding-options hash-key]
family inet {
 layer-3;
 layer-4;
 symmetric-hash {
 [complement;]
```

```

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

```

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

```

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

```

```

 }
 layer-4;
 }
}
symmetric-hash {
 [complement;]
}
}
}
.
.
.
}
}

```

### Configuring Symmetrical Load Balancing on Trio-Based MPCs

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

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

- Compute a Symmetrical Hash

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

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

- Configure Link Indexes

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

- Enable Symmetric Load Balancing

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

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

- Achieve Symmetry for Bridged and Routed Traffic

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


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

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

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

Traffic polarization occurs when LAGs are configured on cascaded routers. For example, in [Figure 4 on page 128](#), if a certain flow uses Link 1 of the aggregated Ethernet bundle between Device R1 and Device R2, the flow also uses Link 1 of the aggregated Ethernet bundle between Device R2 and Device R3.

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



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

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

## Example Configurations

### IN THIS SECTION

- [Example Configurations of Chassis Wide Settings | 129](#)
- [Example Configurations of Per-Packet-Forwarding-Engine Settings | 130](#)

### *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 | 133](#)
- [Configuring Symmetrical Hashing for family inet on Both Routers | 134](#)
- [Configuring Symmetrical Hashing for family inet and family multiservice on the Two Routers | 135](#)

**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 | 136](#)

## Example: Configuring Aggregated Ethernet Load Balancing

### IN THIS SECTION

- [Requirements | 136](#)
- [Overview | 136](#)
- [Configuration | 139](#)
- [Verification | 153](#)

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 | 139](#)

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 ae x 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 ae x aggregated-ether-options load-balance** command.

## Topology

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

Figure 5: Aggregated Ethernet Load Balancing



## Configuration

### IN THIS SECTION

- [CLI Quick Configuration | 140](#)
- [Configuring Adaptive Load Balancing | 145](#)
- [Results | 149](#)

### **CLI Quick Configuration**

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

#### **R1**

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

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

R2

```
set chassis aggregated-devices ethernet device-count 5
set interfaces ge-1/2/0 unit 0 family inet address 120.168.100.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-1/2/1 unit 0 family inet address 120.168.101.1/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-1/3/0 gigether-options 802.3ad ae0
set interfaces ge-1/3/1 gigether-options 802.3ad ae0
set interfaces ge-1/3/2 gigether-options 802.3ad ae0
set interfaces ge-1/3/3 gigether-options 802.3ad ae0
set interfaces ge-1/3/4 gigether-options 802.3ad ae0
set interfaces ge-2/2/1 gigether-options 802.3ad ae1
set interfaces ge-2/2/2 gigether-options 802.3ad ae1
set interfaces ge-2/2/3 gigether-options 802.3ad ae1
set interfaces ge-2/2/4 gigether-options 802.3ad ae1
set interfaces ge-2/2/5 gigether-options 802.3ad ae1
set interfaces ge-2/2/6 gigether-options 802.3ad ae1
set interfaces ge-2/2/7 gigether-options 802.3ad ae1
set interfaces ge-2/2/8 gigether-options 802.3ad ae1
set interfaces ae0 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.1/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options load-balance adaptive tolerance 10
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 unit 0 family inet address 120.168.105.1/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0004.00
set accounting-options selective-aggregate-interface-stats disable
set protocols rsvp interface ge-1/2/0.0
set protocols rsvp interface ge-1/2/1.0
set protocols rsvp interface ae0.0
```

```

set protocols rsvp interface ae1.0
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ge-1/2/0.0
set protocols isis interface ge-1/2/1.0
set protocols isis interface ae0.0
set protocols isis interface ae1.0
set protocols isis interface lo0.0

```

### R3

```

set chassis aggregated-devices ethernet device-count 5
set interfaces xe-4/0/0 unit 0 family inet address 120.168.9.1/30
set interfaces xe-4/0/0 unit 0 family mpls
set interfaces xe-4/0/1 unit 0 family inet address 120.168.10.1/30
set interfaces xe-4/0/1 unit 0 family mpls
set interfaces ge-5/0/1 gigether-options 802.3ad ae1
set interfaces ge-5/0/2 gigether-options 802.3ad ae1
set interfaces ge-5/0/3 gigether-options 802.3ad ae1
set interfaces ge-5/0/4 gigether-options 802.3ad ae1
set interfaces ge-5/0/5 gigether-options 802.3ad ae1
set interfaces ge-5/0/6 gigether-options 802.3ad ae1
set interfaces ge-5/0/7 gigether-options 802.3ad ae1
set interfaces ge-5/0/8 gigether-options 802.3ad ae1
set interfaces ge-5/3/0 gigether-options 802.3ad ae0
set interfaces ge-5/3/1 gigether-options 802.3ad ae0
set interfaces ge-5/3/2 gigether-options 802.3ad ae0
set interfaces ge-5/3/3 gigether-options 802.3ad ae0
set interfaces ge-5/3/4 gigether-options 802.3ad ae0
set interfaces ae0 aggregated-ether-options link-speed 1g
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family inet address 120.168.104.2/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 aggregated-ether-options link-speed 1g
set interfaces ae1 aggregated-ether-options lacp active

```

```
set interfaces ae1 unit 0 family inet address 120.168.105.2/30
set interfaces ae1 unit 0 family iso
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 120.168.0.9/32
set interfaces lo0 unit 0 family iso address 49.0001.1201.6800.0009.00
set routing-options router-id 120.168.0.9
set routing-options autonomous-system 55
set protocols rsvp interface xe-4/0/0.0
set protocols rsvp interface xe-4/0/1.0
set protocols rsvp interface ae0.0
set protocols rsvp interface ae1.0
set protocols mpls label-switched-path to-videl to 120.168.0.2
set protocols mpls interface xe-4/0/0.0
set protocols mpls interface xe-4/0/1.0
set protocols mpls interface ae0.0
set protocols mpls interface ae1.0
set protocols bgp group pe-routers type internal
set protocols bgp group pe-routers local-address 120.168.0.9
set protocols bgp group pe-routers family inet unicast
set protocols bgp group pe-routers family inet-vpn unicast
set protocols bgp group pe-routers neighbor 120.168.0.2
set protocols isis traffic-engineering family inet shortcuts
set protocols isis level 1 disable
set protocols isis interface ae0.0
set protocols isis interface ae1.0
set protocols isis interface lo0.0
set policy-options policy-statement nhs then next-hop self
set policy-options policy-statement vpn-m5-export term 1 from protocol bgp
set policy-options policy-statement vpn-m5-export term 1 from protocol direct
set policy-options policy-statement vpn-m5-export term 1 then community add vpn-m5-target
set policy-options policy-statement vpn-m5-export term 1 then accept
set policy-options policy-statement vpn-m5-export term 2 then reject
set policy-options policy-statement vpn-m5-import term 1 from protocol bgp
set policy-options policy-statement vpn-m5-import term 1 from protocol direct
set policy-options policy-statement vpn-m5-import term 1 from community vpn-m5-target
set policy-options policy-statement vpn-m5-import term 1 then accept
set policy-options policy-statement vpn-m5-import term 2 then reject
set policy-options community vpn-m5-target members target:55:100
set routing-instances vpn-m5 instance-type vrf
set routing-instances vpn-m5 interface xe-4/0/0.0
```

```

set routing-instances vpn-m5 interface xe-4/0/1.0
set routing-instances vpn-m5 route-distinguisher 120.168.0.9:1
set routing-instances vpn-m5 vrf-import vpn-m5-import
set routing-instances vpn-m5 vrf-export vpn-m5-export
set routing-instances vpn-m5 protocols bgp group ce type external
set routing-instances vpn-m5 protocols bgp group ce peer-as 100
set routing-instances vpn-m5 protocols bgp group ce as-override
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.9.2
set routing-instances vpn-m5 protocols bgp group ce neighbor 120.168.10.2
set routing-instances vpn-m5 protocols ospf domain-id 1.0.0.0
set routing-instances vpn-m5 protocols ospf export vpn-m5-import
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/0.0
set routing-instances vpn-m5 protocols ospf area 0.0.0.0 interface xe-4/0/1.0

```

### ***Configuring Adaptive Load Balancing***

#### **Step-by-Step Procedure**

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

To configure the R2 router:

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

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

```

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

```

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

```

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

```

```

user@R2# set ge-1/2/1 unit 0 family iso
user@R2# set ge-1/2/1 unit 0 family mpls
user@R2# set lo0 unit 0 family inet address 120.168.0.4/32
user@R2# set lo0 unit 0 family iso address 49.0001.1201.6800.0004.00

```

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

```

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

```

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

```

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

```

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

```

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

```

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

```

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

```

7. Configure LACP on the ae0 aggregated Ethernet bundle.

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

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

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

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

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

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

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

11. Configure LACP on the ae1 aggregated Ethernet bundle.

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

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

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

13. Disable selective aggregate Ethernet statistics.

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

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

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

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

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

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

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

## Results

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

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

```
user@R2# show interfaces
ge-1/2/0 {
 unit 0 {
 family inet {
 address 120.168.100.1/30;
 }
 family iso;
 family mpls;
 }
}
ge-1/2/1 {
 unit 0 {
 family inet {
 address 120.168.101.1/30;
 }
 family iso;
 family mpls;
 }
}
ge-1/3/0 {
 gigether-options {
 802.3ad ae0;
 }
}
ge-1/3/1 {
 gigether-options {
 802.3ad ae0;
 }
}
```

```
}
ge-1/3/2 {
 gigether-options {
 802.3ad ae0;
 }
}
ge-1/3/3 {
 gigether-options {
 802.3ad ae0;
 }
}
ge-1/3/4 {
 gigether-options {
 802.3ad ae0;
 }
}
ge-2/2/1 {
 gigether-options {
 802.3ad ae1;
 }
}
ge-2/2/2 {
 gigether-options {
 802.3ad ae1;
 }
}
ge-2/2/3 {
 gigether-options {
 802.3ad ae1;
 }
}
ge-2/2/4 {
 gigether-options {
 802.3ad ae1;
 }
}
ge-2/2/5 {
 gigether-options {
 802.3ad ae1;
 }
}
ge-2/2/6 {
 gigether-options {
```

```
 802.3ad ae1;
 }
}
ge-2/2/7 {
 gigether-options {
 802.3ad ae1;
 }
}
ge-2/2/8 {
 gigether-options {
 802.3ad ae1;
 }
}
ae0 {
 aggregated-ether-options {
 load-balance {
 adaptive tolerance 10;
 }
 link-speed lg;
 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 lg;
 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 | 153](#)

Confirm that the configuration is working properly.

### ***Verifying Adaptive Load Balancing on ae0***

#### **Purpose**

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

#### **Action**

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

```

user@R2> show interfaces ae0 extensive
Logical interface ae0.0 (Index 325) (SNMP ifIndex 917) (Generation 134)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
 Statistics Packets pps Bytes bps
Bundle:
 Input : 848761 9 81247024 7616
 Output: 166067308909 3503173 126900990064983 21423804256
Adaptive Statistics:
 Adaptive Adjusts: 264
 Adaptive Scans : 27682
 Adaptive Updates: 10
Link:
 ge-1/3/0.0

```

```

Input : 290888 5 29454436 3072
Output: 33183442699 704569 25358563587277 4306031760
ge-1/3/1.0
Input : 162703 1 14806325 992
Output: 33248375409 705446 25406995966732 4315342152
ge-1/3/2.0
Input : 127448 1 12130566 992
Output: 33184552729 697572 25354827700261 4267192376
ge-1/3/3.0
Input : 121044 1 11481262 1280
Output: 33245875402 697716 25405953405192 4265750584
ge-1/3/4.0
Input : 146678 1 13374435 1280
Output: 33205071207 697870 25374651121458 4269487384

```

## Meaning

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

## SEE ALSO

[Aggregated Ethernet Interfaces](#) | 61

### Release History Table

| 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.                   |

### Release History Table

| 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.                                 |

## RELATED DOCUMENTATION

[Aggregated Ethernet Interfaces | 61](#)

[Link Protection of Aggregated Ethernet Interfaces | 106](#)

[Understanding Ethernet Link Aggregation on ACX Series Routers | 163](#)

## 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 | 155](#)
- [Guidelines for Configuring Performance Monitoring Functionalities on Aggregated Ethernet Interfaces | 158](#)

Use this topic to understand more 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 Release 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 graceful 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 active-standby 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 software-based. 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\*](#)

### Release History Table

| 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 Release 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 | 61](#)

*ITU-T Y.1731 Ethernet Service OAM Overview*

## Periodic Packet Management

### IN THIS SECTION

- [Understanding Periodic Packet Management on MX Series Routers | 160](#)
- [Configuring Periodic Packet Management on MX Series Routers | 160](#)

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

PPM operates in two modes:

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

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

**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

- [Identifying Periodic Packet Management Mode | 161](#)
- [Enabling Centralized Periodic Packet Management | 162](#)

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.

This topic describes:

## Identifying Periodic Packet Management Mode

Before you configure periodic packet management, you must identify the mode of periodic packet management.

To identify the mode of periodic packet management:

1. From operational mode, enter the **show ppm adjacencies detail** command.

```
user@host> show ppm adjacencies detail
Protocol: OSPF2, Hold time: 40000, IFL-index: 359
Distributed: FALSE
OSPF source key: 88.1.1.2, OSPF area ID: 0.0.0.0
```

In the above example, the distributed field is false. So, the periodic packet management mode for the OSPF protocol is centralized or running on the Routing Engine only.

OR

From configuration mode, enter the **run show ppm adjacencies detail** command.

```
user@host# run show ppm adjacencies detail
Protocol: BFD, Hold time: 900, IFL-index: 359
Distributed: TRUE
BFD discriminator: 16, BFD routing table index: 0
```

In the above example, the distributed field is true. So, the periodic packet management mode for the BFD protocol is distributed to PFE.

2. From configuration mode, enter the **run show ppm adjacencies protocol *protocol-name* detail** command.

```
user@host# show ppm adjacencies protocol lacp detail
Protocol: LACP, Hold time: 3000, IFL-index: 361
Distributed: TRUE
Distribution handle: 30, Distribution address: fpc1
Adjacencies: 1, Remote adjacencies: 1
```

In the above example, the distributed field is true. So, the periodic packet management mode for the LACP protocol is distributed to PFE.

**NOTE:** You can also run the `show ppm adjacencies` command from the PFE shell. When you run the command from the PFE shell, the command displays all the process that are running in distributed mode.

## Enabling Centralized Periodic Packet Management

After you identify the periodic packet management mode, you can enable centralized periodic packet management. When you enable centralized periodic packet management, the `ppm` process runs on the routing engine only. When you enable centralized periodic packet management, you have disabled distributed PPM. You can enable centralized periodic packet management for troubleshooting to identify if the protocol is having issues while running on distributed mode. If you do not face the issue while the protocol is running on centralized mode, you can narrow down the issue and identify if the issue is because of 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 `no-delegate-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 **run show ppm adjacencies detail** command.

```
user@host# run show ppm adjacencies detail
 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 \(Ethernet Switching\) | 1082](#)

## Understanding Ethernet Link Aggregation on ACX Series Routers

### IN THIS SECTION

- [Load Balancing | 166](#)
- [LACP Monitoring | 167](#)
- [Link Protection | 168](#)
- [Understanding the Algorithm Used to Hash LAG Bundle | 169](#)

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 (ae $x$ ), that is, the defined bundle, to be labeled "up":

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

```
[edit interfaces]
user@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 overridden 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 13 on page 170](#) 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 13: 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<br>Destination MAC<br>Source MAC and Destination MAC |
| IP           | Source IP and Destination IP  | Source MAC Address<br>Destination MAC<br>Source MAC and Destination MAC |
| MPLS         | MPLS label 1 and MPLS label 2 | Source MAC Address<br>Destination MAC<br>Source MAC and Destination MAC |

[Table 14 on page 171](#) 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 14: 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)<br>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*

# 2

PART

## Gigabit Ethernet Interfaces

---

[Configuring Gigabit Ethernet Interfaces](#) | 173

[Configure Port Speed](#) | 293

[Configuring Optical Transport network](#) | 458

---

# Configuring Gigabit Ethernet Interfaces

## IN THIS CHAPTER

- [Configuring 10-Gigabit Ethernet PICs | 173](#)
- [Configuring 40-Gigabit Ethernet PICs | 215](#)
- [Configuring 100-Gigabit Ethernet MICs/PICs | 218](#)
- [Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | 246](#)
- [Configuring Layer 2 Overhead Attribute in Interface Statistics | 259](#)
- [Configuring Gigabit Ethernet Policers | 268](#)
- [Gigabit Ethernet Autonegotiation | 284](#)

## Configuring 10-Gigabit Ethernet PICs

### IN THIS SECTION

- [10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 174](#)
- [12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 178](#)
- [24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 181](#)
- [P2-10G-40G-QSFPP PIC Overview | 182](#)
- [Configuring the P2-10G-40G-QSFPP PIC | 193](#)
- [Example: Configuring the P2-10G-40G-QSFPP PIC | 197](#)
- [Framing Overview | 202](#)
- [Understanding WAN Framing | 203](#)
- [Configuring Ethernet Framing | 203](#)
- [Modes of Operation | 204](#)
- [Configuring Mixed-Rate Mode Operation | 205](#)
- [Configuring Line-Rate Mode on PICs Supporting Oversubscription | 207](#)

- [Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC | 207](#)
- [Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 209](#)
- [Gigabit Ethernet Notification of Link Down Alarm Overview | 212](#)
- [Notification of Link Down for Optics Options Overview | 213](#)
- [Configuring Gigabit Ethernet Notification of Link Down Alarm | 213](#)
- [Configuring Link Down Notification for Optics Options Alarm or Warning | 214](#)

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 queue 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 15 on page 178](#) 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP).

**Table 15: 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)<br>960 filters across multiple logical interfaces |
| MAC policers            | 128 ingress Mac policers<br>128 egress Mac policers                                                  |
| Classifiers             | 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*

*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.<br>For IPv6 the limit is 1022. |
| Classifiers                         | Eight classifiers per PIC for each BA classifier type      |

## SEE ALSO

| [Configuring 40-Gigabit Ethernet PICs](#)

## 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*.
- 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*

*Gigabit Ethernet Autonegotiation*

## P2-10G-40G-QSFPP PIC Overview

### IN THIS SECTION

- [Understanding Dual Configuration on P2-10G-40G-QSFPP PIC | 183](#)
- [Understanding Port Group | 184](#)
- [Port Numbering on P2-10G-40G-QSFPP PIC When Port Groups Are Not Configured | 188](#)
- [10-Gigabit Ethernet Mode | 191](#)
- [40-Gigabit Ethernet Mode | 193](#)

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 16 on page 185](#) shows the port numbering in 40-Gigabit Ethernet mode and in 10-Gigabit Ethernet mode at the port group level.

**Table 16: Port Number Mapping When Port Groups Are Configured**

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

Table 16: 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<br>et-1/1/1:1<br>et-1/1/1:2<br>et-1/1/1:3 |
|                   | et-1/1/2                                   | et-1/1/2:0<br>et-1/1/2:1<br>et-1/1/2:2<br>et-1/1/2:3 |
| 3(1)              | et-1/1/3                                   | et-1/1/3:0<br>et-1/1/3:1<br>et-1/1/3:2<br>et-1/1/3:3 |
|                   | et-1/1/4                                   | et-1/1/4:0<br>et-1/1/4:1<br>et-1/1/4:2<br>et-1/1/4:3 |
|                   | et-1/1/5                                   | et-1/1/5:0<br>et-1/1/5:1<br>et-1/1/5:2<br>et-1/1/5:3 |

Table 16: 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<br>et-1/1/6:1<br>et-1/1/6:2<br>et-1/1/6:3 |
|                   | et-1/1/7                                   | et-1/1/7:0<br>et-1/1/7:1<br>et-1/1/7:2<br>et-1/1/7:3 |
|                   | et-1/1/8                                   | et-1/1/8:0<br>et-1/1/8:1<br>et-1/1/8:2<br>et-1/1/8:3 |
| 6(2)              | et-1/1/6                                   | et-1/1/6:0<br>et-1/1/6:1<br>et-1/1/6:2<br>et-1/1/6:3 |
|                   | et-1/1/7                                   | et-1/1/7:0<br>et-1/1/7:1<br>et-1/1/7:2<br>et-1/1/7:3 |

Table 16: 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<br>et-1/1/8:1<br>et-1/1/8:2<br>et-1/1/8:3     |
| 9(3)              | et-1/1/9                                   | et-1/1/9:0<br>et-1/1/9:1<br>et-1/1/9:2<br>et-1/1/9:3     |
|                   | et-1/1/10                                  | et-1/1/10:0<br>et-1/1/10:1<br>et-1/1/10:2<br>et-1/1/10:3 |
|                   | et-1/1/11                                  | et-1/1/11:0<br>et-1/1/11:1<br>et-1/1/11:2<br>et-1/1/11:3 |

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

Table 17 on page 189 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 17: 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<br>et-1/1/0:1<br>et-1/1/0:2<br>et-1/1/0:3 |
| 1                 | et-1/1/1                                   | et-1/1/1:0<br>et-1/1/1:1<br>et-1/1/1:2<br>et-1/1/1:3 |
| 2                 | et-1/1/2                                   | et-1/1/2:0<br>et-1/1/2:1<br>et-1/1/2:2<br>et-1/1/2:3 |
| 3                 | et-1/1/3                                   | et-1/1/3:0<br>et-1/1/3:1<br>et-1/1/3:2<br>et-1/1/3:3 |
| 4                 | et-1/1/4                                   | et-1/1/4:0<br>et-1/1/4:1<br>et-1/1/4:2<br>et-1/1/4:3 |

Table 17: 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<br>et-1/1/5:1<br>et-1/1/5:2<br>et-1/1/5:3 |
| 6                 | et-1/1/6                                   | et-1/1/6:0<br>et-1/1/6:1<br>et-1/1/6:2<br>et-1/1/6:3 |
| 7                 | et-1/1/7                                   | et-1/1/7:0<br>et-1/1/7:1<br>et-1/1/7:2<br>et-1/1/7:3 |
| 8                 | et-1/1/8                                   | et-1/1/8:0<br>et-1/1/8:1<br>et-1/1/8:2<br>et-1/1/8:3 |
| 9                 | et-1/1/9                                   | et-1/1/9:0<br>et-1/1/9:1<br>et-1/1/9:2<br>et-1/1/9:3 |

Table 17: 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<br>et-1/1/10:1<br>et-1/1/10:2<br>et-1/1/10:3 |
| 11                | et-1/1/11                                  | et-1/1/11:0<br>et-1/1/11:1<br>et-1/1/11:2<br>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* otn-options preemptive-fast-reroute]** hierarchy level.
- Configure thresholds and interval for the optical channel data unit (ODU) signal degradation (**odu-signal-degrade**) and the configurable pre-FEC bit error rate (BER) (**ber-threshold-signal-degrade**) at the **[edit interfaces *interface-name* otn-options odu-signal-degrade]** hierarchy level and the **[edit interfaces *interface-name* otn-options signal-degrade]** hierarchy level, respectively.

## 40-Gigabit Ethernet Mode

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

### SEE ALSO

[Configuring 100-Gigabit Ethernet MICs/PICs | 218](#)

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

### IN THIS SECTION

- [Configuring the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 193](#)
- [Configuring the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 194](#)
- [Configuring the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 194](#)
- [Configuring the PIC at Port Group Level | 195](#)
- [Configuring Framing Mode on P2-10G-40G-QSFPP PIC | 196](#)

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 | 196](#)
- [Configuring LAN PHY Framing Mode in 40-Gigabit Ethernet Mode | 197](#)

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 Ethernet 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](#) | 182

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

### IN THIS SECTION

- [Requirements](#) | 198
- [Overview](#) | 198
- [Configuration](#) | 198

## 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 | 200](#)

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 | 201](#)

## 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
 Bit errors Seconds
 Errored blocks 1
 Interface transmit statistics: Disabled

```

### Meaning

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

### SEE ALSO

| [Configuring 100-Gigabit Ethernet MICs/PICs](#) | 218

## 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](#)

## Understanding WAN Framing

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

In particular:

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

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

### SEE ALSO

*framing*

*Configuring Ethernet Framing*

*Framing Overview*

[Ethernet Interfaces User Guide for Routing Devices](#)

## Configuring Ethernet Framing

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

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

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

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

**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*

**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 1G |10G** 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 **1G** by including the `set fpc fpc-slot pic pic-number port port-number speed 1G` statement at the `[edit chassis]` 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*  
*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](#)

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

[Table 18 on page 207](#) 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 18: 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 18: 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).<br><br>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*

*Configuring Line-Rate Mode on PICs Supporting Oversubscription*

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

**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 6 on page 210](#).

**Figure 6: 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.

```
[edit 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*.

**SEE ALSO**

| [Configuring Link Down Notification for Optics Options Alarm or Warning](#)

**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 gigheter-options]** hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
 gigheter-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);
 }
 }
}
```

### Release History Table

| 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 | 218](#)

[Configuring Gigabit Ethernet Policers | 268](#)

[Gigabit Ethernet Autonegotiation | 284](#)

## Configuring 40-Gigabit Ethernet PICs

### IN THIS SECTION

- [40-Gigabit Ethernet PIC Overview | 215](#)
- [Configuring 40-Gigabit Ethernet PICs | 217](#)

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*. 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*.

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*, *Configuring Gigabit Ethernet Policers*, and [Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview](#).

## SEE ALSO

*Configuring Gigabit Ethernet Policers*

*Configuring Gigabit Ethernet Policers*

[Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview](#)

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 268](#)

[Gigabit Ethernet Autonegotiation | 284](#)

## Configuring 100-Gigabit Ethernet MICs/PICs

### IN THIS SECTION

● [100-Gigabit Ethernet Interfaces Overview | 219](#)

● [MPC3E MIC Overview | 222](#)

- [100-Gigabit Ethernet Type 4 PIC with CFP Overview | 224](#)
- [Configuring 100-Gigabit Ethernet Type 4 PIC With CFP | 227](#)
- [Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 231](#)
- [100-Gigabit Ethernet Type 5 PIC with CFP Overview | 233](#)
- [100-Gigabit Ethernet Interfaces Interoperability | 235](#)
- [Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | 238](#)
- [Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 239](#)
- [Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | 241](#)
- [Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 242](#)

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 | 219](#)
- [PTX Series 100-Gigabit Ethernet Interfaces | 221](#)
- [T Series 100-Gigabit Ethernet Interfaces | 221](#)

### MX Series 100-Gigabit Ethernet Interfaces

[Table 19 on page 220](#) lists the 100-Gigabit Ethernet interfaces supported by MX Series routers.

Table 19: 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<br>MX480<br>MX960<br>MX2010<br>MX2020 | <a href="#">100-Gigabit Ethernet MIC with CFP</a><br><i>MPC3E MIC Overview</i> |
| 100-Gigabit Ethernet MIC with CXP           | MIC3-3D-1X100GE-CXP | MX240<br>MX480<br>MX960<br>MX2010<br>MX2020 | <a href="#">100-Gigabit Ethernet MIC with CXP</a><br><i>MPC3E MIC Overview</i> |
| 100-Gigabit Ethernet ports on the MPC4E     | MPC4E-3D-2CGE-8XGE  | MX240<br>MX480<br>MX960<br>MX2010<br>MX2020 | <a href="#">2x100GE + 8x10GE MPC4E</a>                                         |
| 100-Gigabit Ethernet MIC with CFP2          | MIC6-100G-CFP2      | MX2010<br>MX2020                            | <a href="#">100-Gigabit Ethernet MIC with CFP2</a>                             |
| 100-Gigabit Ethernet MIC with CXP (4 Ports) | MIC6-100G-CXP       | MX2010<br>MX2020                            | <a href="#">100-Gigabit Ethernet MIC with CXP (4 Ports)</a>                    |

## PTX Series 100-Gigabit Ethernet Interfaces

[Table 20 on page 221](#) lists the 100-Gigabit Ethernet interfaces supported by PTX Series routers.

**Table 20: 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            | <a href="#">100-Gigabit Ethernet PIC with CFP (PTX Series)</a>                                                                                                                    |
| 100-Gigabit Ethernet PIC with CFP2 | P2-100GE-CFP2      | PTX5000            | <a href="#">100-Gigabit Ethernet PIC with CFP2 (PTX Series)</a>                                                                                                                   |
| 100-Gigabit Ethernet OTN PIC       | P2-100GE-OTN       | PTX5000            | <a href="#">100-Gigabit Ethernet OTN PIC with CFP2 (PTX Series)</a><br><br><i>Understanding the P2-100GE-OTN PIC</i><br><br><i>Configuring OTN Interfaces on P2-100GE-OTN PIC</i> |
| 100-Gigabit DWDM OTN PIC           | P1-PTX-2-100G-WDM  | PTX5000<br>PTX3000 | <a href="#">100-Gigabit DWDM OTN PIC (PTX Series)</a>                                                                                                                             |

## T Series 100-Gigabit Ethernet Interfaces

[Table 21 on page 222](#) lists the 100-Gigabit Ethernet interfaces supported by T Series routers.

**Table 21: 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<br>T4000    | <a href="#">100-Gigabit Ethernet PIC with CFP (T1600 Router)</a><br><br><a href="#">100-Gigabit Ethernet PIC with CFP (T4000 Router)</a><br><br><i>100-Gigabit Ethernet Type 4 PIC with CFP Overview</i><br><br><i>Configuring 100-Gigabit Ethernet Type 4 PIC With CFP</i> |
| 100-Gigabit Ethernet PIC with CFP (Type 5) | PF-1CGE-CFP     | T4000             | <a href="#">100-Gigabit Ethernet PIC with CFP (T4000 Router)</a><br><br><i>100-Gigabit Ethernet Type 5 PIC with CFP Overview</i>                                                                                                                                            |

**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-QSFP)

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 QSFP MIC supports the 40BASE-SR4 interface and uses quad small form-factor pluggable (QSFP) 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* 50-Gigabit 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 \(T1600 Router\)](#)

## 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 PIC0 **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/0/0:0** and **et-x/0/0:1** 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*, *Configuring Gigabit Ethernet Policers*, 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*

*Configuring Gigabit Ethernet Policers*

[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 PFE0.
- 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 interoperability 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*, *Configuring Gigabit Ethernet Policers*, 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 22 on page 235 lists the capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP.

**Table 22: 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.<br>For IPv6 the limit is 1022. |

## SEE ALSO

*Configuring Gigabit Ethernet Policers*

*Configuring Gigabit Ethernet Policers*

[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](#) | 236

- [Interoperability of the MPC4E-3D-2CGE-8XGE MPC with PICs on Other Routers | 236](#)
- [Interoperability of the P1-PTX-2-100GE-CFP PIC with PICs on Other Routers | 237](#)
- [Interoperability of the PD-1CE-CFP-FPC4 PIC with PICs or MICs on Other Routers | 237](#)

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 23 on page 236](#) lists the Interoperability with the 100-Gigabit Ethernet MIC with CFP.

**Table 23: 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) | <a href="#">Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode</a> |

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

[Table 24 on page 236](#) lists the Interoperability with the MPC4E.

**Table 24: MPC4E Interoperability**

| Interoperates with... |  | For More Information... |
|-----------------------|--|-------------------------|
|                       |  |                         |

|          |                                                             |                                                                                                                                             |
|----------|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| T Series | 100-Gigabit Ethernet PIC with CFP (Type 4) (PD-1CECFP-FPC4) | <a href="#">Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode</a> |
|----------|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|

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

[Table 25 on page 237](#) lists the Interoperability with 100-Gigabit Ethernet PIC with CFP (Type 5).

**Table 25: 100-Gigabit Ethernet PIC with CFP (Type 5) (P1-PTX-2-100GE-CFP) Interoperability**

| Interoperates with... |                                                               | For More Information...                                                                                                                                                                                                                 |
|-----------------------|---------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| T Series              | 100-Gigabit Ethernet PIC with CFP (Type 4) (PD- 1CE-CFP-FPC4) | <p><i>Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP</i></p> <p><i>Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4</i></p> |

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

[Table 26 on page 237](#) lists the 100-Gigabit Ethernet PIC with CFP (Type 4).

**Table 26: 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) | <p><i>Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4</i></p> <p><i>forwarding-mode</i></p> <p><i>sa-multicast</i></p> |
| MX Series             | 100-Gigabit Ethernet MIC with CFP (MIC3-3D-1X100GE-CFP)  | <a href="#">Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode</a>                         |

|            |                                                                 |                                                                                                                                                                                                                                         |
|------------|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|            | 100-Gigabit Ethernet ports on the MPC4E                         | <a href="#">Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode</a>                                                                                             |
| PTX Series | 100-Gigabit Ethernet PIC with CFP (Type 5) (P1-PTX-2-100GE-CFP) | <p><i>Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP</i></p> <p><i>Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4</i></p> |

## SEE ALSO

[Periodic Packet Management | 159](#)

[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 Gbps 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 | 239](#)

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-Gbps 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 | 243](#)
- [Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface | 244](#)

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 {
 port 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 together-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](#)

### Release History Table

| 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](#) | 268

| [Gigabit Ethernet Autonegotiation](#) | 284

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

### IN THIS SECTION

- [Transporting Legacy Traffic over Packet Switched Networks](#) | 246
- [Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH Traffic](#) | 249
- [Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy SDH Traffic](#) | 254

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](#) | 247
- [Smart SFP Transceivers for Transporting SDH Traffic over PSNs Overview](#) | 248

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 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.

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).

**NOTE:** The MPC4E (MPC4E-3D-32XGE-SFP 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 | 249](#)
- [Overview | 250](#)
- [Configuring the DS3 Smart SFP | 250](#)
- [Verification | 252](#)

### 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 | 250](#)

## 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
```

- 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
```

- (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
```

- (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
```

- (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 | 253](#)

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 0
 Rx unicast frames 1187124
 Rx multicast frames 2
 Rx broadcast frames 0
 Rx fragments 0
 Rx undersize frames 0
 Rx oversize frames 0
 Rx invalid vlan mismatch frames 0
 Tx frames 1392780998
 Tx bytes 1796396824
 Tx unicast frames 1377974

```

```

Tx multicast frames 1391403024
Tx broadcast frames 0
Smart SFP AV IWF Encap/Decap Statistics: Counters
Rx Packets 0
Tx Packets 0
Malformed Packets 0
Reordered Packets 0
Misordered Dropped Packets 0
Missing Packets 0
PlayedOut Packets 0
JitterBuffer Overrun 0
JitterBuffer Underrun 0
Smart SFP DS3 port[P0] statistics: Counters
BiPolarVariations/Excessive zero errors 0
Tx B3
Errors 0
Code Violation path errors 0
Logical interface ge-4/1/0.0 (Index 350) (SNMP ifIndex 615)
Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
Input packets : 1384454023
Output packets: 0
Protocol multiservice, MTU: Unlimited

```

## Meaning

The DS3 Smart SFP is configured on the MX480 router and you can view the DS3 Smart SFP statistics.

## 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 | 255](#)
- [Overview | 255](#)
- [Configuring the STM1 Smart SFP | 255](#)
- [Verification | 257](#)

## 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 | 255](#)

### *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 | 258](#)

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 0
 Rx unicast frames 1187124
 Rx multicast frames 2
 Rx broadcast frames 0
 Rx fragments 0
 Rx undersize frames 0
 Rx oversize frames 0
 Rx invalid vlan mismatch frames 0
 Tx frames 1392780998
 Tx bytes 1796396824
 Tx unicast frames 1377974

```

```

Tx multicast frames 1391403024
Tx broadcast frames 0
Smart SFP AV IWF Encap/Decap Statistics: Counters
Rx Packets 0
Tx Packets 0
Malformed Packets 0
Reordered Packets 0
Misordered Dropped Packets 0
Missing Packets 0
PlayedOut Packets 0
JitterBuffer Overrun 0
JitterBuffer Underrun 0
Smart SFP STM1 port[P0] statistics: Counters
 BiPolarVariations/Excessive zero errors 0
 Tx B3
Errors 0
 Code Violation path errors 0
Logical interface ge-3/1/0.0 (Index 350) (SNMP ifIndex 615)
 Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
 Input packets : 1384454023
Output packets: 0
 Protocol multiservice, MTU: Unlimited

```

## Meaning

The STM1 Smart SFP is configured on the MX480 router and you can view the STM1 Smart SFP statistics.

## Configuring Layer 2 Overhead Attribute in Interface Statistics

### IN THIS SECTION

- [Accounting of the Layer 2 Overhead Attribute in Interface Statistics | 260](#)
- [Configuring Layer 2 Overhead Accounting in Interface Statistics | 263](#)
- [Verifying the Accounting of Layer 2 Overhead in Interface Statistics | 264](#)

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 | 262](#)

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 27 on page 261 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 27: 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                          |
| 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 | 263](#)

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 | 264](#)
- [Action | 264](#)

### 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: 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
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
 L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
 Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped
packets
 0 best-effort 882558
882558 0
 1 expedited-fo 0
0 0
 2 assured-forw 0
0 0
 3 network-cont 3232
3232 0
Active alarms : None
Active defects : None

```

Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)

Flags: SNMP-Traps 0x4000 Encapsulation: ENET2

Egress account overhead: 100

Ingress account overhead: 90

Traffic statistics:

```

Input bytes : 271524
Output bytes : 37769598
Input packets: 3664
Output packets: 885790

```

IPv6 transit statistics:

```

Input bytes : 0
Output bytes : 16681118
Input packets: 0
Output packets: 362633

```

Local statistics:

```

Input bytes : 271524
Output bytes : 308560
Input packets: 3664
Output packets: 3659

```

Transit statistics:

```

Input bytes : 0 0 bps
Output bytes : 37461038 0 bps
Input packets: 0 0 pps
Output packets: 882131 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 16681118
Input packets: 0
Output packets: 362633
Multicast statistics:
IPV4 multicast statistics:
Input bytes : 112048 0 bps
Output bytes : 20779920 0 bps
Input packets: 1801 0 pps
Output packets: 519498 0 pps
IPV6 multicast statistics:
Input bytes : 156500 0 bps
Output bytes : 16681118 0 bps
Input packets: 1818 0 pps
Output packets: 362633 0 pps
Protocol inet, MTU: 1500, Generation: 151, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 40.40.40.0/30, Local: 40.40.40.2, Broadcast: 40.40.40.3,
Generation: 167
Protocol inet6, MTU: 1500, Generation: 152, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: ::40.40.40.0/126, Local: ::40.40.40.2
Generation: 169
Addresses, Flags: Is-Preferred
Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:d974
Protocol multiservice, MTU: Unlimited, Generation: 171
Generation: 153, Route table: 0
Policer: Input: __default_arp_policer__

```

**SEE ALSO**


---

*show interfaces*

*show interfaces statistics*

### Release History Table

| 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](#)

[Understanding Optical Transport Network \(OTN\) | 458](#)

## Configuring Gigabit Ethernet Policers

### IN THIS SECTION

- [Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 268](#)
- [Configuring Gigabit Ethernet Policers | 271](#)
- [Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 279](#)

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 tri-color 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 28 on page 269](#) 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 28: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs**

| Capability                           | Gigabit Ethernet IQ (SFP) | Gigabit Ethernet (SFP) |
|--------------------------------------|---------------------------|------------------------|
| <b>Layer 2</b>                       |                           |                        |
| 802.3ad link aggregation             | Yes                       | Yes                    |
| Maximum VLANs per port               | 384                       | 1023                   |
| Maximum transmission unit (MTU) size | 9192                      | 9192                   |
| MAC learning                         | Yes                       | Yes                    |
| MAC accounting                       | Yes                       | Yes                    |
| MAC filtering                        | Yes                       | Yes                    |
| Destinations per port                | 960                       | 960                    |
| Sources per port                     | 64                        | 64                     |

Table 28: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs (Continued)

| Capability                                                              | Gigabit Ethernet IQ (SFP)  | Gigabit Ethernet (SFP) |
|-------------------------------------------------------------------------|----------------------------|------------------------|
| 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                     |
| <b>Layer 2 VPNs</b>                                                     |                            |                        |
| VLAN CCC                                                                | Yes                        | Yes                    |
| Port-based CCC                                                          | Yes                        | Yes                    |
| Extended VLAN CCC Virtual Metropolitan Area Network (VMAN) Tag Protocol | Yes                        | Yes                    |
| <b>CoS</b>                                                              |                            |                        |
| PIC-based egress queues                                                 | Yes                        | Yes                    |
| Queued VLANs                                                            | Yes                        | No                     |
| VPLS                                                                    | Yes                        | Yes                    |

For more information about configuring VPLS, see the [Junos OS VPNs Library for Routing Devices](#).

You can also configure CoS on logical IQ interfaces. For more information, see the [Junos OS Class of Service User Guide for Routing Devices](#).

## SEE ALSO

*[Configuring MAC Address Accounting](#)*

*[Configuring a Policer Overhead](#)*

## Configuring Gigabit Ethernet Policers

### IN THIS SECTION

- [Overview | 271](#)
- [Configuring a Policer | 272](#)
- [Specifying an Input Priority Map | 273](#)
- [Specifying an Output Priority Map | 273](#)
- [Applying a Policer | 274](#)
- [Configuring MAC Address Filtering | 276](#)
- [Example: Configuring Gigabit Ethernet Policers | 277](#)

### 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 policers allow you to perform simple traffic policing without configuring a firewall filter. First you configure the Ethernet policer profile, next you classify ingress and egress traffic, then you can apply the policer to a logical interface.

For Gigabit Ethernet PICs with SFPs (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* ggether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces interface-name ggether-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* gigeather-options ethernet-policer-profile input-priority-map]** hierarchy level:

```
[edit interfaces interface-name gigeather-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"](#).

**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 gigeother-options ethernet-policer-profile output-priority-map]` hierarchy level:

```
[edit interfaces interface-name gigeother-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 29 on page 274](#) shows the predefined forwarding classes and their associated queue assignments.

**Table 29: 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](#)".

### 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 policers allow you to perform simple traffic policing without configuring a firewall filter.

To apply policers to specific source MAC addresses, include the **accept-source-mac** statement:

```
accept-source-mac {
 mac-address mac-address {
 policer {
 input cos-policer-name;
 output cos-policer-name;
 }
 }
}
```

You can include these statements at the following hierarchy levels:

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

You can specify the MAC address as *nn:nn:nn:nn:nn:nn* or *nnnn.nnnn.nnnn*, where *n* is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include multiple **mac-address** statements in the logical interface configuration.

**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](#)".

To enable this blocking, include the **source-filtering** statement at the **[edit interfaces *interface-name* ggether-options]** hierarchy level:

```
[edit interfaces interface-name ggether-options]
source-filtering;
```

For more information about the **source-filtering** statement, see *Configuring MAC Address Filtering for Ethernet Interfaces*.

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* gigheter-options]** hierarchy level:

```
[edit interfaces interface-name gigheter-options]
no-source-filtering;
```

## Example: Configuring Gigabit Ethernet Policers

### IN THIS SECTION

- [Example | 277](#)
- [Example Configuration | 277](#)

### *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 {
 gigheter-options {
 ether-switch-profile {
 ether-policer-profile {
 input-priority-map {
 ieee-802.1p {
```



```
}
}
```

## RELATED DOCUMENTATION

*Configuring MAC Address Accounting*

*Configuring a Policer Overhead*

## Configuring Gigabit Ethernet Two-Color and Tricolor Policers

### IN THIS SECTION

- [Overview | 279](#)
- [Configuring a Policer | 280](#)
- [Applying a Policer | 281](#)
- [Example: Configuring and Applying a Policer | 282](#)

### 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*), 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 [edit firewall] hierarchy level. For each interface, you can configure a three-color policer or two-color input policer or output policers—you cannot configure both a three-color policer and a two-color policer.

### 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](#)

[Configuring a Policer Overhead](#)

## Release History Table

| 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 | 173](#)

[Configuring 40-Gigabit Ethernet PICs | 215](#)

## Gigabit Ethernet Autonegotiation

### IN THIS SECTION

- [Gigabit Ethernet Autonegotiation Overview | 284](#)
- [Configuring Gigabit Ethernet Autonegotiation | 284](#)

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](#)

### Configuring Gigabit Ethernet Autonegotiation

#### IN THIS SECTION

- [Configuring Gigabit Ethernet Autonegotiation with Remote Fault | 285](#)
- [Configuring Flow Control | 285](#)

- [Configuring Autonegotiation Speed on MX Series Routers | 285](#)
- [Displaying Autonegotiation Status | 286](#)

## 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 gige-opts]` hierarchy level.

```
[edit interfaces ge-fpc/pic/port gige-opts]
(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 gige-opts]` hierarchy level. For more information, see *Configuring Flow Control*.

## 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, the **auto-10m-100m** option allows the fixed tri-speed port to auto negotiate with ports limited by **100m** or **10m** maximum speed. This option must be enabled only for Tri-rate MPC port, that is, 3D 40x 1GE (LAN) RJ45 MIC 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-fpc/pic/port extensive**.

[Table 30 on page 286](#) and [Table 31 on page 290](#) 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 30: 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 |                        |

Table 30: Mode and Autonegotiation Status (Local) (Continued)

| Transmit | Receive | Mode               | LED   | Link | Autonegotiation Status |
|----------|---------|--------------------|-------|------|------------------------|
| 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 |                        |
| 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 |                        |

Table 30: Mode and Autonegotiation Status (Local) (Continued)

| Transmit | Receive | Mode                 | LED   | Link | Autonegotiation Status |
|----------|---------|----------------------|-------|------|------------------------|
| 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     |
| 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               |

Table 30: Mode and Autonegotiation Status (Local) (Continued)

| Transmit | Receive | Mode                 | LED   | Link | Autonegotiation Status |
|----------|---------|----------------------|-------|------|------------------------|
| 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 |                        |
| 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 |                        |

**Table 30: Mode and Autonegotiation Status (Local) (Continued)**

| Transmit | Receive | Mode                | LED   | Link | Autonegotiation Status |
|----------|---------|---------------------|-------|------|------------------------|
| ON       | ON      | Explicit+RFI-Online | Green | UP   | Complete               |

**Table 31: Mode and Autonegotiation Status (Remote)**

| Transmit | Receive | Mode               | LED   | Link | Autonegotiation Status |
|----------|---------|--------------------|-------|------|------------------------|
| ON       | ON      | Default            | Green | UP   | Complete               |
| ON       | ON      | Default            | Red   | DOWN |                        |
| ON       | OFF     | Default            | Red   | DOWN |                        |
| OFF      | ON      | Default            | Red   | DOWN |                        |
| OFF      | OFF     | Default            | Red   | DOWN |                        |
| ON       | ON      | No-autonegotiation | Green | UP   | Incomplete             |
| ON       | ON      | No-autonegotiation | Red   | DOWN |                        |
| ON       | OFF     | No-autonegotiation | Red   | DOWN |                        |
| OFF      | ON      | No-autonegotiation | Green | UP   |                        |
| OFF      | OFF     | No-autonegotiation | Red   | DOWN |                        |
| ON       | ON      | Explicit           | Green | UP   | Complete               |
| ON       | ON      | Explicit           | Red   | DOWN |                        |

Table 31: Mode and Autonegotiation Status (Remote) *(Continued)*

| Transmit | Receive | Mode                 | LED   | Link | Autonegotiation Status |
|----------|---------|----------------------|-------|------|------------------------|
| 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 |                        |
| 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 |                        |

**Table 31: Mode and Autonegotiation Status (Remote) (Continued)**

| Transmit | Receive | Mode                | LED | Link | Autonegotiation Status |
|----------|---------|---------------------|-----|------|------------------------|
| OFF      | OFF     | Explicit+RFI-Online | Red | DOWN |                        |

**RELATED DOCUMENTATION**

*Configuring 10-Gigabit Ethernet PICs*

*Configuring Layer 2 Overhead Attribute in Interface Statistics*

**Release History Table**

| Release | Description                                                                                                                                                                                                                                                                                                           |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 14.2    | Starting from Junos OS Release 14.2, the <code>auto-10m-100m</code> option allows the fixed tri-speed port to auto negotiate with ports limited by 100m or 10m maximum 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 <code>auto-10m-100m</code> 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 | 215](#)

[Configuring 100-Gigabit Ethernet MICs/PICs | 218](#)

# Configure Port Speed

## IN THIS CHAPTER

- Port Speed | 293
- Configuring the Port Speed on the JNP10K-LC1201 by Using New Port Profile Configuration | 452

## Port Speed

### SUMMARY

Learn about port speed on a device or line card, support for multiple port speed details, guidelines and how to configure the port speed.

### IN THIS SECTION

- Port Speed Overview | 294
- Port Speed on MX10003 MPC Overview | 306
- Port Speed on MPC10E-10C-MRATE Overview | 312
- Port Speed on MPC10E-15C-MRATE Overview | 317
- Port Speed on MX2K-MPC11E Overview | 322
- Port Speed on PTX10K-LC1201 Overview | 334
- Port Speed on ACX7100-48L Router Overview | 345
- Interface Naming Conventions for ACX7100-48L | 347
- Port Speed on ACX710 Router Overview | 348

- [Port Speed on PTX10001-36MR Router Overview | 350](#)
- [Port Speed on PTX10K-LC1202-36MR Overview | 361](#)
- [Introduction to Port Speed | 370](#)
- [Interface Naming Conventions | 394](#)
- [Configuring Port Speed on Routing Devices and Line Cards | 418](#)

## Port Speed Overview

### IN THIS SECTION

- [Interface Naming Conventions | 304](#)
- [What is Oversubscription? | 305](#)

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 32 on page 295](#) describes the different types of port speed configuration.

Table 32: 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 <b>pic-mode</b> statement, then the port operates in the default speed.                                                                                                                                                                                                                                                                                                     |
| Port Level                    | <p>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.</p> <p>When you change the speed of a specific port in a given PIC using the <b>speed</b> 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.</p> |

Table 33: Port Speed Configuration at PIC Level- Chassis Hierarchy

| Configuration Steps                                                | PIC Level                                                                                                                                                                                           |
|--------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1: At the PIC level, specify the operating speed for the PIC. | <pre>[edit chassis fpc <i>fpc-slot</i> pic <i>pic-number</i>] user@host# set pic-mode <i>pic-speed</i></pre> <p>For example:</p> <pre>[edit chassis fpc 0 pic 0] user@host# set pic-mode 100g</pre> |

Step 2: (Optional) At the PIC level, configure the number of active physical ports that operate at the speed configured in Step 2.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set number-of-ports number-of-active-physical-ports
```

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set number-of-ports 4
```

Step 3: (Optional) Specify the number of logical interfaces that you want to create on a physical port.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set number-of-sub-ports number-of-sub-ports
```

For example:

```
[edit chassis fpc 0 pic 1]
user@host# set number-of-sub-ports 4
```

Step 4: (Optional) Specify the port number that you want to power off.

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

For example:

```
[edit chassis fpc 0 pic 1]
user@host# set port 2 Unused
```

Step 5: Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
pic-mode 100G;
number-of-ports 4;
```

```
[edit chassis fpc 0 pic 1]
user@host# show
number-of-sub-ports 4;
```

Step 6: Commit the configuration.

**Table 34: Port Speed Configuration at Port Level– Chassis Hierarchy**

| Configuration Steps                                                                                            | Port Level                                                                                                                                                                                                                                                                          |
|----------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 1: At the port level, specify the operating speed for the port.</p>                                    | <pre>[edit chassis fpc <i>fpc-slot</i> pic <i>pic-number</i>] user@host# <b>set port <i>port-number</i> speed (10g   40g   100g)</b></pre> <p>For example:</p> <pre>[edit chassis fpc 0 pic 0] user@host# <b>set port 0 speed 40g</b> user@host# <b>set port 1 speed 100g</b></pre> |
| <p>Step 2: (Optional) Specify the number of logical interfaces that you want to create on a physical port.</p> | <pre>[edit chassis fpc <i>fpc-slot</i> pic <i>pic-number</i>] user@host# <b>set number-of-sub-ports <i>number-of-sub-ports</i></b></pre> <p>For example:</p> <pre>[edit chassis fpc 0 pic 1] user@host# <b>set number-of-sub-ports 4</b></pre>                                      |
| <p>Step 3: (Optional) Specify the port number that you want to power off.</p>                                  | <pre>[edit chassis fpc <i>fpc-slot</i> pic <i>pic-number</i>] user@host# <b>set port <i>port-number</i> unused</b></pre> <p>For example:</p> <pre>[edit chassis fpc 0 pic 1] user@host# <b>set port 2 Unused</b></pre>                                                              |

**Table 34: Port Speed Configuration at Port Level- Chassis Hierarchy (Continued)**

| Configuration Steps                      | Port Level                                                                                                                                                                                   |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 4: Verify the configuration.</p> | <pre>[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; }</pre> |
| <p>Step 6: Commit the configuration.</p> |                                                                                                                                                                                              |

Table 35 on page 298 describes the steps to configure the port speed for non-channelized interfaces from the [edit interfaces] hierarchy.

**Table 35: Port Speed Configuration for non-channelized interfaces -Interfaces Hierarchy**

| Configuration Steps                                                                                                               | Non-Channelized Interfaces                                                                                                                                                                              |
|-----------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 1: To indicate the speed at which the ports operate, configure the <b>speed</b> statement for the desired interfaces.</p> | <pre>[edit interfaces interface-name] user@host# <b>set speed (10g   25g   40g   50g  100g   400g)</b></pre> <p>For example:</p> <pre>[edit interfaces et-1/0/3] user@host# <b>set speed 100g</b></pre> |

Table 35: Port Speed Configuration for non-channelized interfaces -Interfaces Hierarchy (Continued)

| Configuration Steps                                                                                                                                                                                                                                                 | Non-Channelized Interfaces                                                                                                                                                                                                                                                                   |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 2: To configure the speed for a group of ports.</p>                                                                                                                                                                                                         | <pre>[edit ] user@host# wildcard range set interfaces interface-name speed speed</pre> <p>For example:</p> <pre>[edit ] user@host# wildcard range set interfaces et-1/0/[0-5] speed 100g</pre>                                                                                               |
| <p>Step 3: To specify the number of interfaces you want to configure per port.</p>                                                                                                                                                                                  | <p>Not applicable</p>                                                                                                                                                                                                                                                                        |
| <p>Step 4: (Optional) To control the number of interfaces created on a physical port, use the <b>unused</b> statement. If you configure a port as unused, no interfaces are created for that port irrespective of the port profile configuration for that port.</p> | <pre>[edit] user@host# set interfaces interface-name unused</pre> <p>For example:</p> <pre>[edit] user@host# set interfaces et-2/0/3 unused</pre> <p>In this example, no interfaces (channelized or non-channelized) are created on port 3 of the line card installed in the FPC slot 2.</p> |

**Table 35: Port Speed Configuration for non-channelized interfaces -Interfaces Hierarchy (Continued)**

| Configuration Steps               | Non-Channelized Interfaces                                                                                                            |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Step 5: Verify the configuration. | <pre> et-x/y/z { speed 100g;      unit 0 {         ...     }     ...      unit N {         ...     } } ...  et-x/y/z { unused; </pre> |
| Step 6: Commit the configuration. |                                                                                                                                       |

Table 36: Port Speed Configuration for channelized interfaces -Interfaces Hierarchy

| Configuration Steps                                                                                                               | Channelized Interfaces                                                                                                                                                                          |
|-----------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 1: To indicate the speed at which the ports operate, configure the <b>speed</b> statement for the desired interfaces.</p> | <pre>[edit interfaces interface-name] user@host# set speed (10g   25g   40g   50g   100g   400g)</pre> <p>For example:</p> <pre>[edit interfaces et-1/0/3] user@host# set speed 100g</pre>      |
| <p>Step 2: To configure the speed for a group of ports.</p>                                                                       | <pre>[edit ] user@host# wildcard range set interfaces interface-name speed speed</pre> <p>For example:</p> <pre>[edit ] user@host# wildcard range set interfaces et-1/0/[7-12] speed 100g</pre> |

Table 36: Port Speed Configuration for channelized interfaces -Interfaces Hierarchy (*Continued*)

| Configuration Steps                                                                                                                                                                                                                                                 | Channelized Interfaces                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Step 3: To specify the number of interfaces you want to configure per port.</p>                                                                                                                                                                                  | <pre data-bbox="841 394 1419 506">[edit interfaces <i>interface-name</i>]     user@host# <b>set number- of-sub-ports <i>number-of-sub-ports</i></b></pre> <p data-bbox="841 590 992 621">For example:</p> <pre data-bbox="841 680 1419 791">[edit interfaces <i>et-1/0/3</i>]     user@host# <b>set number- of-sub-ports 4</b></pre> <p data-bbox="841 875 1341 947">In this example, in Step 1 and Step 2, you configure 4x100GE channelized interfaces.</p> |
| <p>Step 4: (Optional) To control the number of interfaces created on a physical port, use the <b>unused</b> statement. If you configure a port as unused, no interfaces are created for that port irrespective of the port profile configuration for that port.</p> | <pre data-bbox="841 1041 1386 1152">[edit]     user@host# <b>set interfaces <i>interface-name</i> unused</b></pre> <p data-bbox="841 1236 992 1268">For example:</p> <pre data-bbox="841 1327 1386 1438">[edit]     user@host# <b>set interfaces <i>et-2/0/4</i> unused</b></pre> <p data-bbox="841 1522 1419 1633">In this example, no interfaces (channelized or non-channelized) are created on port 4 of the line card installed in the FPC slot 2.</p>   |

Table 36: Port Speed Configuration for channelized interfaces -Interfaces Hierarchy (*Continued*)

| Configuration Steps               | Channelized Interfaces                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 5: Verify the configuration. | <pre> et-x/y/z {     speed     100g;     number- of-sub-ports 4;      z:0 {         et-x/y/         unit 0{             ...         }     }      z:1 {         et-x/y/         unit 0{             ...         }     }      z:2 {         et-x/y/         unit 0{             ...         }     }      z:3 {         et-x/y/         unit 0{             ...         }     }     ... </pre> |

**Table 36: Port Speed Configuration for channelized interfaces -Interfaces Hierarchy (Continued)**

| Configuration Steps | Channelized Interfaces                                          |
|---------------------|-----------------------------------------------------------------|
|                     | <pre>et-x/y/z:6 {                                 unused;</pre> |

Step 6: Commit the configuration.

### 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 37 on page 305](#) describes the naming formats for the channelized and non-channelized interfaces.

**Table 37: Channelized and Non-Channelized Interface Naming Formats**

| Interfaces                                                                                                                                                             | Non-channelized Interfaces Naming Formats                                               | Channelized Interfaces Naming Formats                                                           |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| 10-Gigabit Ethernet Interfaces                                                                                                                                         | Prefix is <b>xe-</b> . The interface name appears in the <b>xe-fpc/pic/port</b> format. | Prefix is <b>xe-</b> . The interface name appears in the <b>xe-fpc/pic/port:channel</b> format. |
| 25-Gigabit Ethernet Interfaces, 40-Gigabit Ethernet Interfaces, 100-Gigabit Ethernet Interfaces, 200-Gigabit Ethernet Interfaces, and 400-Gigabit Ethernet Interfaces. | Prefix is <b>et-</b> . The interface name appears in the <b>et-fpc/pic/port</b> format. | Prefix is <b>et-</b> . The interface name appears in the <b>et-fpc/pic/port:channel</b> 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.

**SEE ALSO**


---

[number-of-ports | 1040](#)


---

[number-of-sub-ports | 1042](#)


---

[pic-mode | 1071](#)


---

[speed | 1134](#)


---

[Unused | 1220](#)

## Port Speed on MX10003 MPC Overview

### IN THIS SECTION

- [Port Speed Support on MX10003 MPC | 308](#)
- [Interface Naming Conventions | 309](#)
- [Supported Active Physical Ports on MX10003 MPC to Prevent Oversubscription | 310](#)

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 38 on page 306](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 38: Port Speed for MX10003 MPC**

| PIC                    | Port Number | Port Speed Supported                                                                                                                                           |
|------------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (Fixed-port PIC) | 0-5         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>All the 4x10-Gigabit Ethernet ports can operate in 1 Gbps mode.                                                |
| PIC 1 (Multi-Rate MIC) | 0-11        | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>On the non-MASCEC MIC, all the 4x10-Gigabit Ethernet ports can operate in 1 Gbps mode. |

Follow these guidelines when you configure the speed of a port:

- By default, the MX10003 router comes up with the PIC mode where all the interface operates at the same speed of 10-Gbps. That is, by default, both the PICs (PIC 0 and PIC 1) operate at 10-Gbps

speed. To use different port speeds, you must first switch to the port mode and then change the default speed.

To change the default speed, you must select a port and configure a different port speed on it and reset both the PICs for the configuration to take effect. For example, select 40GE or 100GE on PIC 1 and 10GE on PIC 0. For this configuration to take effect, you must reset both PICs.

- Regardless of the line card— MIC (PIC1) or fixed-port PIC (PIC0) installed —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.
- The port speed configuration on the fixed-port PIC and the MIC must be homogenous. However, at port level you can configure port speeds in heterogeneous mode.

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.

- When you configure the port speed at the port level, only the configured ports are active. Other ports are disabled.
- 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 the ports which 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.
- You cannot configure the speed of the port at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations.
- When you change an existing port speed configuration at the port level, you must reset the PIC for the configuration to take effect. When you change an existing port speed configuration at the PIC level, the MPC automatically resets the PIC.

For information on how to configure the speed at the PIC level, see [Table 33 on page 295](#). For information on how to configure the speed at the port level, see [Table 34 on page 297](#).

## Port Speed Support on MX10003 MPC

Table 39 on page 308 summarizes the port mode configuration at the Packet Forwarding Engine level.

**Table 39: PFE Based Port Mode Configuration**

| Port Speed configuration on PIC1(Gbps) |       |       |       | Port speed configuration on PIC0(Gbps) |       |
|----------------------------------------|-------|-------|-------|----------------------------------------|-------|
| 100                                    | 100   | 100   | 100   | 0                                      | 0     |
| 100                                    | 100   | 100   | 10/40 | 0                                      | 0     |
| 100                                    | 100   | 100   | 0     | 10/40                                  | 10/40 |
| 100                                    | 100   | 10/40 | 10/40 | 10/40                                  | 10/40 |
| 100                                    | 10/40 | 10/40 | 10/40 | 10/40                                  | 10/40 |
| 10/40                                  | 10/40 | 10/40 | 0     | 10/40                                  | 10/40 |
| 10/40                                  | 10/40 | 10/40 | 10/40 | 0                                      | 0     |

**NOTE:** Use the [port-checker tool](#) to check whether the combination of ports you want to use is valid or not.

Table 40 on page 308 summarizes the PIC mode configuration.

**Table 40: PIC Mode Configuration**

| Port Speed configuration on PIC1(Gbps) | Port speed configuration on PIC0(Gbps)   |
|----------------------------------------|------------------------------------------|
| 100                                    | 0<br>Configure the number of ports to 0. |

**Table 40: PIC Mode Configuration (Continued)**

| Port Speed configuration on PIC1(Gbps) | Port speed configuration on PIC0(Gbps) |
|----------------------------------------|----------------------------------------|
| 10                                     | 10                                     |
| 40                                     | 40                                     |

### Interface Naming Conventions

[Table 41 on page 309](#) lists the naming conventions used for interfaces on the fixed-port PIC when installed in slot 0 of the MX10003 MPC. [Table 42 on page 310](#) lists the naming conventions used for interfaces on the modular MIC when installed in slot 1 of the MPC.

**Table 41: 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                        | <i>xe-x/0/0:[0-3]</i>         | <i>et-x/0/0</i>               | -                              |
|                          | <i>xe-x/0/1:[0-3]</i>         | <i>et-x/0/1</i>               | -                              |
| 1                        | <i>xe-x/0/2:[0-3]</i>         | <i>et-x/0/2</i>               | -                              |
|                          | <i>xe-x/0/3:[0-3]</i>         | <i>et-x/0/3</i>               | -                              |
| 2                        | <i>xe-x/0/4:[0-3]</i>         | <i>et-x/0/4</i>               | -                              |
|                          | <i>xe-x/0/5:[0-3]</i>         | <i>et-x/0/5</i>               | -                              |

**Table 42: 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                        | <i>xe-x/1/0:[0-3]</i>         | <b>et-x/1/0</b>               | <b>et-x/1/0</b>                |
|                          | <i>xe-x/1/1:[0-3]</i>         | <b>et-x/1/1</b>               | <b>et-x/1/1</b>                |
|                          | <i>xe-x/1/2:[0-3]</i>         | <b>et-x/1/2</b>               | <b>et-x/1/2</b>                |
|                          | <i>xe-x/1/3:[0-3]</i>         | <b>et-x/1/3</b>               | <b>et-x/1/3</b>                |
| 1                        | <i>xe-x/1/4:[0-3]</i>         | <b>et-x/1/4</b>               | <b>et-x/1/4</b>                |
|                          | <i>xe-x/1/5:[0-3]</i>         | <b>et-x/1/5</b>               | <b>et-x/1/5</b>                |
|                          | <i>xe-x/1/6:[0-3]</i>         | <b>et-x/1/6</b>               | <b>et-x/1/6</b>                |
|                          | <i>xe-x/1/7:[0-3]</i>         | <b>et-x/1/7</b>               | <b>et-x/1/7</b>                |
| 2                        | <i>xe-x/1/8:[0-3]</i>         | <b>et-x/1/8</b>               | <b>et-x/1/8</b>                |
|                          | <i>xe-x/1/9:[0-3]</i>         | <b>et-x/1/9</b>               | <b>et-x/1/9</b>                |
|                          | <i>xe-x/1/10:[0-3]</i>        | <b>et-x/1/10</b>              | <b>et-x/1/10</b>               |
|                          | <i>xe-x/1/11:[0-3]</i>        | <b>et-x/1/11</b>              | <b>et-x/1/11</b>               |

**Supported Active Physical Ports on MX10003 MPC to Prevent Oversubscription**

[Table 43 on page 311](#) lists the active physical ports on MX10003 MPC at the MIC Level.

Table 43: 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               |
| 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 44 on page 312 list the active physical ports on MX10003 MPC at PIC Level.

**Table 44: Active Physical Ports on MX10003 MPC at the PIC level**

| Ports Configured (number-of-ports Statement) | Active Physical Ports for Different Configured Speeds |                     |
|----------------------------------------------|-------------------------------------------------------|---------------------|
|                                              | 10-Gigabit Ethernet                                   | 40-Gigabit Ethernet |
| 1                                            | 0                                                     | 0                   |
| 2                                            | 0, 2                                                  | 0, 2                |
| 3                                            | 0, 2, 4                                               | 0, 2, 4             |
| 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 MPC10E-10C-MRATE Overview

### IN THIS SECTION

- [Port Speed Support on MPC10E-10C-MRATE | 314](#)
- [Interface Naming Conventions | 314](#)
- [Supported Active Physical Ports on MPC10E-10C-MRATE to Prevent Oversubscription | 315](#)

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

Table 45 on page 313 summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 45: Port Speed for MPC10E-10C-MRATE**

| PIC              | Port Number | Port Speed Supported                                                                                                                                                                          |
|------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (or PFE 0) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet (Only on port 4)<br>Default: All active ports operate in 100 Gbps mode. |
| PIC 1 (or PFE 1) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet (Only on port 4)<br>Default: All active ports operate in 100 Gbps mode. |

Follow these guidelines when you configure the speed of a port:

- By default, the MPC10E-10C-MRATE comes up with the PIC mode where all the interface operates at the same speed of 100-Gbps. That is, by default, all the PICs (PIC 0 and PIC 1) operate at 100-Gbps speed.
- When you change the speed at the PIC level, the PIC restarts automatically with the new configured speed.
- When you configure the speed of a port at the port level, then only the speed of that port is modified. All other ports in the PIC remain unaffected.
- You cannot configure the port speed at the PIC level and the port level simultaneously. Commit fails when you configure the port speed at the PIC level and port level simultaneously.

For information on how to configure the speed at the PIC level, see [Table 33 on page 295](#). For information on how to configure the speed at the port level, see [Table 34 on page 297](#).

### Port Speed Support on MPC10E-10C-MRATE

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 46 on page 314](#) summarizes the port profile configuration on MPC10E-10C-MRATE.

**Table 46: 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                |

### Interface Naming Conventions

[Table 47 on page 314](#) lists the interface naming conventions for the MPC10E-10C-MRATE.

**Table 47: 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   | <i>xe-x/0/0:[0-3]</i>         | <i>et-x/0/0:[0-3]</i>         | <i>et-x/0/0</i>               | <i>et-x/0/0</i>                | -                              |
|     | <i>xe-x/0/1:[0-3]</i>         | <i>et-x/0/1:[0-3]</i>         | <i>et-x/0/1</i>               | <i>et-x/0/1</i>                | -                              |
|     | <i>xe-x/0/2:[0-3]</i>         | <i>et-x/0/2:[0-3]</i>         | <i>et-x/0/2</i>               | <i>et-x/0/2</i>                | -                              |

Table 47: 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 |
|-----|-------------------------------|-------------------------------|-------------------------------|--------------------------------|--------------------------------|
|     | <i>xe-x/0/3:[0-3]</i>         | <i>et-x/0/3:[0-3]</i>         | <i>et-x/0/3</i>               | <i>et-x/0/3</i>                | -                              |
|     | <i>xe-x/0/4:[0-3]</i>         | <i>et-x/0/4:[0-3]</i>         | <i>et-x/0/4</i>               | <i>et-x/0/4</i>                | <i>et-x/0/4</i>                |
| 1   | <i>xe-x/1/0:[0-3]</i>         | <i>et-x/1/0:[0-3]</i>         | <i>et-x/1/0</i>               | <i>et-x/1/0</i>                | -                              |
|     | <i>xe-x/1/1:[0-3]</i>         | <i>et-x/1/1:[0-3]</i>         | <i>et-x/1/1</i>               | <i>et-x/1/1</i>                | -                              |
|     | <i>xe-x/1/2:[0-3]</i>         | <i>et-x/1/2:[0-3]</i>         | <i>et-x/1/2</i>               | <i>et-x/1/2</i>                | -                              |
|     | <i>xe-x/1/3:[0-3]</i>         | <i>et-x/1/3:[0-3]</i>         | <i>et-x/1/3</i>               | <i>et-x/1/3</i>                | -                              |
|     | <i>xe-x/1/4:[0-3]</i>         | <i>et-x/1/4:[0-3]</i>         | <i>et-x/1/4</i>               | <i>et-x/1/4</i>                | <i>et-x/0/4</i>                |

#### Supported Active Physical Ports on MPC10E-10C-MRATE to Prevent Oversubscription

Table 48 on page 316 lists the active ports with port speed configuration at PIC level for MPC10E-10C-MRATE.

Table 48: Active Ports with port speed configured at PIC level

| PIC Type              | Number of Active Ports | 10-Gigabit Ethernet | 25-Gigabit Ethernet | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
|-----------------------|------------------------|---------------------|---------------------|---------------------|----------------------|----------------------|
| 5xQSF P28 PIC (PIC 0) | 1                      | 0                   | 0                   | 0                   | 0                    | 4                    |
|                       | 2                      | 0, 1                | 0,1                 | 0, 1                | 0, 1                 | 4                    |
|                       | 3                      | 0, 1, 2             | 0,1,2               | 0, 1, 2             | 0, 1, 2              | 4                    |
|                       | 4                      | 0, 1, 2, 3          | 0,1,2,3             | 0, 1, 2, 3          | 0, 1, 2, 3           | 4                    |
|                       | 5                      | 0, 1, 2, 3, 4       | 0,1,2,3,4           | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |
| 5xQSF P28 PIC (PIC 1) | 1                      | 0                   | 0                   | 0                   | 0                    | 4                    |
|                       | 2                      | 0, 1                | 0,1                 | 0, 1                | 0, 1                 | 4                    |
|                       | 3                      | 0, 1, 2             | 0,1,2               | 0, 1, 2             | 0, 1, 2              | 4                    |
|                       | 4                      | 0, 1, 2, 3          | 0,1,2,3             | 0, 1, 2, 3          | 0, 1, 2, 3           | 4                    |
|                       | 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 MPC10E-15C-MRATE Overview

### IN THIS SECTION

- [Port Speed Support on MPC10E-15C-MRATE | 319](#)
- [Interface Naming Conventions | 319](#)
- [Supported Active Physical Ports on MPC10E-15C-MRATE to Prevent Oversubscription | 321](#)

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 49 on page 317](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 49: Port Speed for MPC10E-15C-MRATE**

| PIC              | Port Number | Port Speed Supported                                                                                                                                                                          |
|------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (or PFE 0) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet (Only on port 4)<br>Default: All active ports operate in 100 Gbps mode. |

Table 49: Port Speed for MPC10E-15C-MRATE (Continued)

| PIC              | Port Number | Port Speed Supported                                                                                                                                                                          |
|------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 1 (or PFE 1) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet (Only on Port 4)<br>Default: All active ports operate in 100 Gbps mode. |
| PIC 2 (or PFE 2) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet (Only on port 4)<br>Default: All active ports operate in 100 Gbps mode. |

Follow these guidelines when you configure the speed of a port:

- By default, the MPC10E-15C-MRATE comes up with the PIC mode where all the interface operates at the same speed of 100-Gbps. That is, by default, all the PICs (PIC 0, PIC1 and PIC 2) operate at 100-Gbps speed.
- When you change the speed at the PIC level, the PIC restarts automatically with the new configured speed.
- When you configure the speed of a port at the port level, then only the speed of that port is modified. All other ports in the PIC remain unaffected.
- You cannot configure the port speed at the PIC level and the port level simultaneously. Commit fails when you configure the port speed at the PIC level and port level simultaneously.

For information on how to configure the speed at the PIC level, see [Table 33 on page 295](#). For information on how to configure the speed at the port level, see [Table 34 on page 297](#).

## Port Speed Support on MPC10E-15C-MRATE

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 50 on page 319](#) summarizes the port profile configuration on MPC10E-15C-MRATE.

**Table 50: Port speed support on MPC10E-15C-MRATE**

| PIC                  | PIC level Profile |     |     |      |      | Port level Profile |
|----------------------|-------------------|-----|-----|------|------|--------------------|
|                      | 10G               | 25G | 40G | 100G | 400G |                    |
| PIC 0 (5xQSFP28 PIC) | Yes               | Yes | Yes | Yes  | Yes  | Yes                |
| PIC 1 (5xQSFP28 PIC) | Yes               | Yes | Yes | Yes  | Yes  | Yes                |
| PIC 2 (5xQSFP28 PIC) | Yes               | Yes | Yes | Yes  | Yes  | Yes                |

## Interface Naming Conventions

[Table 51 on page 319](#) lists the interface naming conventions for the MPC10E-15C-MRATE.

**Table 51: 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   | <b><i>xe-x/0/0:[0-3]</i></b>  | <b><i>et-x/0/0:[0-3]</i></b>  | <b><i>et-x/0/0</i></b>        | <b><i>et-x/0/0</i></b>         | -                              |
|     | <b><i>xe-x/0/1:[0-3]</i></b>  | <b><i>et-x/0/1:[0-3]</i></b>  | <b><i>et-x/0/1</i></b>        | <b><i>et-x/0/1</i></b>         | -                              |
|     | <b><i>xe-x/0/2:[0-3]</i></b>  | <b><i>et-x/0/2:[0-3]</i></b>  | <b><i>et-x/0/2</i></b>        | <b><i>et-x/0/2</i></b>         | -                              |

Table 51: 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 |
|-----|-------------------------------|-------------------------------|-------------------------------|--------------------------------|--------------------------------|
|     | <i>xe-x/0/3:[0-3]</i>         | <i>et-x/0/3:[0-3]</i>         | <i>et-x/0/3</i>               | <i>et-x/0/3</i>                | -                              |
|     | <i>xe-x/0/4:[0-3]</i>         | <i>et-x/0/4:[0-3]</i>         | <i>et-x/0/4</i>               | <i>et-x/0/4</i>                | <i>et-x/0/4</i>                |
| 1   | <i>xe-x/1/0:[0-3]</i>         | <i>et-x/1/0:[0-3]</i>         | <i>et-x/1/0</i>               | <i>et-x/1/0</i>                | -                              |
|     | <i>xe-x/1/1:[0-3]</i>         | <i>et-x/1/1:[0-3]</i>         | <i>et-x/1/1</i>               | <i>et-x/1/1</i>                | -                              |
|     | <i>xe-x/1/2:[0-3]</i>         | <i>et-x/1/2:[0-3]</i>         | <i>et-x/1/2</i>               | <i>et-x/1/2</i>                | -                              |
|     | <i>xe-x/1/3:[0-3]</i>         | <i>et-x/1/3:[0-3]</i>         | <i>et-x/1/3</i>               | <i>et-x/1/3</i>                | -                              |
|     | <i>xe-x/1/4:[0-3]</i>         | <i>et-x/1/4:[0-3]</i>         | <i>et-x/1/4</i>               | <i>et-x/1/4</i>                | <i>et-x/0/4</i>                |
| 2   | <i>xe-x/2/0:[0-3]</i>         | <i>et-x/2/0:[0-3]</i>         | <i>et-x/2/0</i>               | <i>et-x/2/0</i>                | -                              |
|     | <i>xe-x/2/1:[0-3]</i>         | <i>et-x/2/1:[0-3]</i>         | <i>et-x/2/1</i>               | <i>et-x/2/1</i>                | -                              |
|     | <i>xe-x/2/2:[0-3]</i>         | <i>et-x/2/2:[0-3]</i>         | <i>et-x/2/2</i>               | <i>et-x/2/2</i>                | -                              |
|     | <i>xe-x/2/3:[0-3]</i>         | <i>et-x/2/3:[0-3]</i>         | <i>et-x/2/3</i>               | <i>et-x/2/3</i>                | -                              |

**Table 51: 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 |
|-----|-------------------------------|-------------------------------|-------------------------------|--------------------------------|--------------------------------|
|     | <b>xe-x/2/4:[0-3]</b>         | <b>et-x/2/4:[0-3]</b>         | <b>et-x/2/4</b>               | <b>et-x/2/4</b>                | <b>et-x/0/4</b>                |

**Supported Active Physical Ports on MPC10E-15C-MRATE to Prevent Oversubscription**

[Table 52 on page 321](#) list the active ports with port speed configuration at PIC level for MPC10E-15C-MRATE.

**Table 52: Active Ports with port speed configuration at PIC level**

| PIC Type                    | Number of Active Ports | 10-Gigabit Ethernet | 25-Gigabit Ethernet | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
|-----------------------------|------------------------|---------------------|---------------------|---------------------|----------------------|----------------------|
| 5xQSFP<br>28 PIC<br>(PIC 0) | 1                      | 0                   | 0                   | 0                   | 0                    | 4                    |
|                             | 2                      | 0, 1                | 0,1                 | 0, 1                | 0, 1                 | 4                    |
|                             | 3                      | 0, 1, 2             | 0,1,2               | 0, 1, 2             | 0, 1, 2              | 4                    |
|                             | 4                      | 0, 1, 2, 3          | 0,1,2,3             | 0, 1, 2, 3          | 0, 1, 2, 3           | 4                    |
|                             | 5                      | 0, 1, 2, 3, 4       | 0,1,2,3,4           | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |
| 5xQSFP<br>28 PIC<br>(PIC 1) | 1                      | 0                   | 0                   | 0                   | 0                    | 4                    |
|                             | 2                      | 0, 1                | 0,1                 | 0, 1                | 0, 1                 | 4                    |

Table 52: Active Ports with port speed configuration at PIC level *(Continued)*

| PIC Type                    | Number of Active Ports | 10-Gigabit Ethernet | 25-Gigabit Ethernet | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
|-----------------------------|------------------------|---------------------|---------------------|---------------------|----------------------|----------------------|
|                             | 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<br>28 PIC<br>(PIC 2) | 1                      | 0                   | 0                   | 0                   | 0                    | 4                    |
|                             | 2                      | 0, 1                | 0,1                 | 0, 1                | 0, 1                 | 4                    |
|                             | 3                      | 0, 1, 2             | 0,1,2               | 0, 1, 2             | 0, 1, 2              | 4                    |
|                             | 4                      | 0, 1, 2, 3          | 0,1,2,3             | 0, 1, 2, 3          | 0, 1, 2, 3           | 4                    |
|                             | 5                      | 0, 1, 2, 3, 4       | 0,1,2,3,4           | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |

## Port Speed on MX2K-MPC11E Overview

### IN THIS SECTION

- [Interface Naming Conventions for MPC11E | 327](#)
- [Supported Active Physical Ports on MX2K-MPC11E to Prevent Oversubscription | 330](#)

For information on the line card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

Table 53 on page 323 summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 53: Port Speed for the MX2K-MPC11E**

| PIC              | Port Number | Port Speed Supported                                                                                                                                                                          |
|------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (or PFE 0) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br><br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br><br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |
| PIC 1 (or PFE 1) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br><br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br><br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |

Table 53: Port Speed for the MX2K-MPC11E (Continued)

| PIC              | Port Number | Port Speed Supported                                                                                                                                                                      |
|------------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 2 (or PFE 2) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |
| PIC 3 (or PFE 3) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |

Table 53: Port Speed for the MX2K-MPC11E (Continued)

| PIC              | Port Number | Port Speed Supported                                                                                                                                                                      |
|------------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 4 (or PFE 4) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |
| PIC 5 (or PFE 5) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |

Table 53: Port Speed for the MX2K-MPC11E (Continued)

| PIC              | Port Number | Port Speed Supported                                                                                                                                                                      |
|------------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 6 (or PFE 6) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |
| PIC 7 (or PFE 7) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode. |
|                  | 1-4         | 100-Gigabit Ethernet<br>Default: All active ports operate in 100-Gigabit Ethernet mode.                                                                                                   |

Follow these guidelines when you configure the speed of a port:

- If you do not configure the port speed at the PIC level using the **pic-mode** option, then the default port speed is 100 Gbps.
- Ports 0 through 4 on each PIC of the MX2K-MPC11E MPC supports speeds of 100 Gbps.

- When you change the speed at the PIC level, the PIC restarts automatically with the new configured speed.
- You cannot configure the port speed at the PIC level and the port level simultaneously. Error messages are displayed when you attempt to commit the configuration.
- When you change the port configuration at the port level, the interfaces corresponding to the affected port are deleted and then re-created.
- You can configure a maximum bandwidth of 500 Gbps per PIC. For example, if you configure the speed of Port 0 of a PIC to 400 Gbps, and the speed of Port 1 of a PIC to 100 Gbps, then you cannot configure any additional ports on that PIC. If you attempt to configure an additional port, it is an invalid configuration.
- If you configure the port speed of a port on any PIC to 100 Gbps and then configure the port speed of port 0 of the same PIC to 400 Gbps, an alarm is generated to indicate that you must restart the PIC. The 400-Gigabit Ethernet interface is created only after you manually restart the PIC.
- When you change an existing port speed configuration with an invalid port speed configuration, an alarm is generated indicating that the port speed configuration is invalid. The MPC continues to operate using the existing port speed configuration or the default port speed.

For information on how to configure the speed at the PIC level, see [Table 33 on page 295](#). For information on how to configure the speed at the port level, see [Table 34 on page 297](#).

### Interface Naming Conventions for MPC11E

[Table 54 on page 327](#) lists the interface naming conventions for the MX2K-MPC11E.

**Table 54: 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   | <i>xe-x/0/0:<br/>[0-3]</i>    | <i>et-x/0/0</i>               | <i>et-x/0/0</i>                | <i>et-x/0/0:[0-3]</i>            | <i>et-x/0/0</i>                |
|     |                               |                               | <i>et-x/0/1</i>                |                                  | -                              |
|     |                               |                               | <i>et-x/0/2</i>                |                                  | -                              |

Table 54: 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/0/3                       |                                  | -                              |
|     |                               |                               | et-x/0/4                       |                                  | -                              |
| 1   | xe-x/1/0:<br>[0-3]            | et-x/1/0                      | et-x/1/0                       | et-x/1/0:[0-3]                   | et-x/1/0                       |
|     |                               |                               | et-x/1/1                       |                                  | -                              |
|     |                               |                               | et-x/1/2                       |                                  | -                              |
|     |                               |                               | et-x/1/3                       |                                  | -                              |
|     |                               |                               | et-x/1/4                       |                                  | -                              |
| 2   | xe-x/2/0:<br>[0-3]            | et-x/2/0                      | et-x/2/0                       | et-x/2/0:[0-3]                   | et-x/2/0                       |
|     |                               |                               | et-x/2/1                       |                                  | -                              |
|     |                               |                               | et-x/2/2                       |                                  | -                              |
|     |                               |                               | et-x/2/3                       |                                  | -                              |
|     |                               |                               | et-x/2/4                       |                                  | -                              |
| 3   | xe-x/3/0:<br>[0-3]            | et-x/3/0                      | et-x/3/0                       | et-x/3/0:[0-3]                   | et-x/3/0                       |
|     |                               |                               | et-x/3/1                       |                                  | -                              |

Table 54: 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/3/2                       |                                  | -                              |
|     |                               |                               | et-x/3/3                       |                                  | -                              |
|     |                               |                               | et-x/3/4                       |                                  | -                              |
| 4   | xe-x/4/0:<br>[0-3]            | et-x/4/0                      | et-x/4/0                       | et-x/4/0:[0-3]                   | et-x/4/0                       |
|     |                               |                               | et-x/4/1                       |                                  | -                              |
|     |                               |                               | et-x/4/2                       |                                  | -                              |
|     |                               |                               | et-x/4/3                       |                                  | -                              |
|     |                               |                               | et-x/4/4                       |                                  | -                              |
| 5   | xe-x/5/0:<br>[0-3]            | et-x/5/0                      | et-x/5/0                       | et-x/5/0:[0-3]                   | et-x/5/0                       |
|     |                               |                               | et-x/5/1                       |                                  | -                              |
|     |                               |                               | et-x/5/2                       |                                  | -                              |
|     |                               |                               | et-x/5/3                       |                                  | -                              |
|     |                               |                               | et-x/5/4                       |                                  | -                              |
| 6   | xe-x/6/0:<br>[0-3]            | et-x/6/0                      | et-x/6/0                       | et-x/6/0:[0-3]                   | et-x/6/0                       |

Table 54: 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/6/1                       |                                  | -                              |
|     |                               |                               | et-x/6/2                       |                                  | -                              |
|     |                               |                               | et-x/6/3                       |                                  | -                              |
|     |                               |                               | et-x/6/4                       |                                  | -                              |
| 7   | xe-x/7/0:<br>[0-3]            | et-x/7/0                      | et-x/7/0                       | et-x/7/0:[0-3]                   | et-x/7/0                       |
|     |                               |                               | et-x/7/1                       |                                  | -                              |
|     |                               |                               | et-x/7/2                       |                                  | -                              |
|     |                               |                               | et-x/7/3                       |                                  | -                              |
|     |                               |                               | et-x/7/4                       |                                  | -                              |

#### Supported Active Physical Ports on MX2K-MPC11E to Prevent Oversubscription

Table 55 on page 331 list the active ports with port speed configuration at PIC level for MX2K-MPC11E.

Table 55: Active Ports with port speed configuration at PIC level

| PIC Type             | Number of Active Ports | 10-Gigabit Ethernet | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
|----------------------|------------------------|---------------------|---------------------|----------------------|----------------------|
| 5xQSFP28 PIC (PIC 0) | 1                      | 0                   | 0                   | 0                    | 0                    |
|                      | 2                      | 0                   | 0                   | 0, 1                 | 0                    |
|                      | 3                      | 0                   | 0                   | 0, 1, 2              | 0                    |
|                      | 4                      | 0                   | 0                   | 0, 1, 2, 3           | 0                    |
|                      | 5                      | 0                   | 0                   | 0, 1, 2, 3, 4        | 0                    |
| 5xQSFP28 PIC (PIC 1) | 1                      | 0                   | 0                   | 0                    | 0                    |
|                      | 2                      | 0                   | 0                   | 0, 1                 | 0                    |
|                      | 3                      | 0                   | 0                   | 0, 1, 2              | 0                    |
|                      | 4                      | 0                   | 0                   | 0, 1, 2, 3           | 0                    |
|                      | 5                      | 0                   | 0                   | 0, 1, 2, 3, 4        | 0                    |
| 5xQSFP28 PIC (PIC 2) | 1                      | 0                   | 0                   | 0                    | 0                    |
|                      | 2                      | 0                   | 0                   | 0, 1                 | 0                    |
|                      | 3                      | 0                   | 0                   | 0, 1, 2              | 0                    |
|                      | 4                      | 0                   | 0                   | 0, 1, 2, 3           | 0                    |

Table 55: 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 |
|----------------------|------------------------|---------------------|---------------------|----------------------|----------------------|
|                      | 5                      | 0                   | 0                   | 0, 1, 2, 3, 4        | 0                    |
| 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                    |

Table 55: 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 |
|----------------------|------------------------|---------------------|---------------------|----------------------|----------------------|
|                      | 4                      | 0                   | 0                   | 0, 1, 2, 3           | 0                    |
|                      | 5                      | 0                   | 0                   | 0, 1, 2, 3, 4        | 0                    |
| 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 PTX10K-LC1201 Overview

### IN THIS SECTION

- [Interface Naming Conventions for PTX10K-LC1201 | 337](#)

For information on the line card, see [PTX10K-LC1201-36D for PTX10008 Routers](#).

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 56 on page 334](#) summarizes the line card details and their description.

**Table 56: PTX10K-LC1201 Software Support and Description**

| Software Feature Support             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forward Error Correction (FEC)       | <ul style="list-style-type: none"> <li>• By default, 200-Gigabit Ethernet interfaces and 400-Gigabit Ethernet interfaces support KP4 FEC mode. You cannot disable FEC mode explicitly.</li> <li>• 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.</li> </ul> <p>For information about FEC support, see "<a href="#">fec (gigether)</a>" on <a href="#">page 713</a>.</p> |
| Maximum Transmission Unit (MTU) Size | Maximum MTU of size 16000 bytes for transit traffic. Ingress interface limits is 9000 bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 56: PTX10K-LC1201 Software Support and Description (Continued)

| Software Feature Support | Description                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Speed Configuration | <ul style="list-style-type: none"> <li>For Junos OS Evolved Releases 19.4R1S1 and 20.1R1, you can configure the port speed at the <b>[edit chassis]</b> hierarchy.</li> <li>For Junos OS Evolved Releases 20.1R2, 20.2R1, and later, you can configure the port speed at the <b>[edit interfaces]</b> hierarchy.</li> </ul> |

Table 57 on page 335 summarizes the Packet Forwarding Engine mapping and the supported port speeds.

Table 57: Port Speed for PTX10K-LC1201

| PIC   | Port Number | Port Speeds Supported                                                                                                                                                                                                                                                                                                                 |
|-------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 | 0-35        | 1x10-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>1x40-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>8x25-Gigabit Ethernet<br>2x50-Gigabit Ethernet<br>1x100-Gigabit Ethernet<br>2x100-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>1x400-Gigabit Ethernet<br><br>Default: All the active ports operate in 400-Gigabit Ethernet mode. |

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 now configure the port speed on the PTX10K-LC1201 line card by using the port profile configuration commands in the

[**edit interfaces**] hierarchy. To streamline the configuration, the new port profile configuration commands are migrated from the [**edit chassis**] hierarchy to the [**edit interfaces**] hierarchy for the PTX10K-LC1201 line card.

Starting in Junos OS Evolved Release 20.4R1, you can configure 10Gbps speed and select the wavelength with SFP+ optics by plugging in the QSA adapter on the QSFP/QSFP+ ports of PTX10K-LC1201 line card. You can now configure the 10Gbps speed by using the **set interfaces *et-**<fpc>/<pic>/<port>** speed <10g>*** or **set interfaces *et-**<fpc>/<pic>/<port>** speed <10g> number-of-sub-ports (1)*** commands. You can select the wavelength of the SFP+ ports by using the **set interfaces *et-**<fpc>/<pic>/<port>**:channel number optics-options wavelength *wavelength**** command and view the output using the **show interfaces diagnostics optics *name*** command.

Follow these guidelines when you configure the speed of a port:

- 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.
- 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.
- You can configure the active physical ports in a PIC. By default, all the physical ports in an interface are active. To control the number of interfaces created on a physical port in a PIC, use the **unused** command. If you configure a port as unused, no interfaces (channelized or non-channelized) are created for that port.
- 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.
- If you have not configured the port profile for a specific port and the port is active, interfaces are created for the port with default speed based on the platform or FPC. You can control the number of interfaces created by using the **unused** command.
- Interfaces are created 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 about how to configure the speed of a port for a non-channelized interface using the new interfaces hierarchy, see [Table 35 on page 298](#). For information about how to configure the speed of a port for a channelized interface using the new interfaces hierarchy, see [Table 36 on page 301](#).

### Interface Naming Conventions for PTX10K-LC1201

[Table 58 on page 337](#) lists the naming conventions used for interfaces on PTX10K-LC1201 for PTX10008 routers. PTX10008 routers support 8 PTX10K-LC1201 line cards.

**Table 58: Interface Naming Convention for PTX10K-LC1201 line card**

| 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 | <i>et-x/0/0:<br/>[0-7]</i>    | <i>et-x/0/0:<br/>[0-7]</i>    | <i>et-x/0/0</i>               | <i>et-x/0/0:<br/>[0-8]</i>    | <i>et-x/0/0<br/>et-x/0/0:[0-1]<br/>et-x/0/0:[0-3]</i> | <i>et-x/0/0:<br/>[0-2]</i>     | <i>et-x/0/0</i>                |
|       | <i>et-x/0/1:<br/>[0-7]</i>    | <i>et-x/0/1:<br/>[0-7]</i>    | <i>et-x/0/1</i>               | <i>et-x/0/1:<br/>[0-8]</i>    | <i>et-x/0/1<br/>et-x/0/1:[0-1]<br/>et-x/0/1:[0-3]</i> | <i>et-x/0/1:<br/>[0-2]</i>     | <i>et-x/0/1</i>                |
|       | <i>et-x/0/2:<br/>[0-7]</i>    | <i>et-x/0/2:<br/>[0-7]</i>    | <i>et-x/0/2</i>               | <i>et-x/0/2:<br/>[0-8]</i>    | <i>et-x/0/2<br/>et-x/0/2:[0-1]<br/>et-x/0/2:[0-3]</i> | <i>et-x/0/2:<br/>[0-2]</i>     | <i>et-x/0/2</i>                |
|       | <i>et-x/0/3:<br/>[0-7]</i>    | <i>et-x/0/3:<br/>[0-7]</i>    | <i>et-x/0/3</i>               | <i>et-x/0/3:<br/>[0-8]</i>    | <i>et-x/0/3<br/>et-x/0/3:[0-1]<br/>et-x/0/3:[0-3]</i> | <i>et-x/0/3:<br/>[0-2]</i>     | <i>et-x/0/3</i>                |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/4:<br>[0-7]            | et-x/0/4:<br>[0-7]            | et-x/0/4                      | et-x/0/4:<br>[0-8]            | et-x/0/4<br>et-x/0/4:[0-1]<br>et-x/0/4:[0-3] | et-x/0/4:<br>[0-2]             | et-x/0/4                       |
|     | et-x/0/5:<br>[0-7]            | et-x/0/5:<br>[0-7]            | et-x/0/5                      | et-x/0/5:<br>[0-8]            | et-x/0/5<br>et-x/0/5:[0-1]<br>et-x/0/5:[0-3] | et-x/0/5:<br>[0-2]             | et-x/0/5                       |
|     | et-x/0/6:<br>[0-7]            | et-x/0/6:<br>[0-7]            | et-x/0/6                      | et-x/0/6:<br>[0-8]            | et-x/0/6<br>et-x/0/6:[0-1]<br>et-x/0/6:[0-3] | et-x/0/6:<br>[0-2]             | et-x/0/6                       |
|     | et-x/0/7:<br>[0-7]            | et-x/0/7:<br>[0-7]            | et-x/0/7                      | et-x/0/7:<br>[0-8]            | et-x/0/7<br>et-x/0/7:[0-1]<br>et-x/0/7:[0-3] | et-x/0/7:<br>[0-2]             | et-x/0/7                       |
|     | et-x/0/8:<br>[0-7]            | et-x/0/8:<br>[0-7]            | et-x/0/8                      | et-x/0/8:<br>[0-8]            | et-x/0/8<br>et-x/0/8:[0-1]<br>et-x/0/8:[0-3] | et-x/0/8:<br>[0-2]             | et-x/0/8                       |
|     | et-x/0/9:<br>[0-7]            | et-x/0/9:<br>[0-7]            | et-x/0/9                      | et-x/0/9:<br>[0-8]            | et-x/0/9<br>et-x/0/9:[0-1]<br>et-x/0/9:[0-3] | et-x/0/9:<br>[0-2]             | et-x/0/9                       |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/10:<br>[0-7]           | et-x/0/10:<br>[0-7]           | et-x/0/10                     | et-x/0/10:<br>[0-8]           | et-x/0/10<br>et-x/0/10:<br>[0-1]<br>et-x/0/10:<br>[0-3] | et-x/0/10:<br>[0-2]            | et-x/0/10                      |
|     | et-x/0/11:<br>[0-7]           | et-x/0/11:<br>[0-7]           | et-x/0/11                     | et-x/0/11:<br>[0-8]           | et-x/0/11<br>et-x/0/11:<br>[0-1]<br>et-x/0/11:<br>[0-3] | et-x/0/11:<br>[0-2]            | et-x/0/11                      |
|     | et-x/0/12:<br>[0-7]           | et-x/0/12:<br>[0-7]           | et-x/0/12                     | et-x/0/12:<br>[0-8]           | et-x/0/12<br>et-x/0/12:<br>[0-1]<br>et-x/0/12:<br>[0-3] | et-x/0/12:<br>[0-2]            | et-x/0/12                      |
|     | et-x/0/13:<br>[0-7]           | et-x/0/13:<br>[0-7]           | et-x/0/13                     | et-x/0/13:<br>[0-8]           | et-x/0/13<br>et-x/0/13:<br>[0-1]<br>et-x/0/13:<br>[0-3] | et-x/0/13:<br>[0-2]            | et-x/0/13                      |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/14:<br>[0-7]           | et-x/0/14:<br>[0-7]           | et-x/0/14                     | et-x/0/14:<br>[0-8]           | et-x/0/14<br>et-x/0/14:<br>[0-1]<br>et-x/0/14:<br>[0-3] | et-x/0/14:<br>[0-2]            | et-x/0/14                      |
|     | et-x/0/15:<br>[0-7]           | et-x/0/15:<br>[0-7]           | et-x/0/15                     | et-x/0/15:<br>[0-8]           | et-x/0/15<br>et-x/0/15:<br>[0-1]<br>et-x/0/15:<br>[0-3] | et-x/0/15:<br>[0-2]            | et-x/0/15                      |
|     | et-x/0/16:<br>[0-7]           | et-x/0/16:<br>[0-7]           | et-x/0/16                     | et-x/0/16:<br>[0-8]           | et-x/0/16<br>et-x/0/16:<br>[0-1]<br>et-x/0/16:<br>[0-3] | et-x/0/16:<br>[0-2]            | et-x/0/16                      |
|     | et-x/0/17:<br>[0-7]           | et-x/0/17:<br>[0-7]           | et-x/0/17                     | et-x/0/17:<br>[0-8]           | et-x/0/17<br>et-x/0/17:<br>[0-1]<br>et-x/0/17:<br>[0-3] | et-x/0/17:<br>[0-2]            | et-x/0/17                      |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/18:<br>[0-7]           | et-x/0/18:<br>[0-7]           | et-x/0/18                     | et-x/0/18:<br>[0-8]           | et-x/0/18<br>et-x/0/18:<br>[0-1]<br>et-x/0/18:<br>[0-3] | et-x/0/18:<br>[0-2]            | et-x/0/18                      |
|     | et-x/0/19:<br>[0-7]           | et-x/0/19:<br>[0-7]           | et-x/0/19                     | et-x/0/19:<br>[0-8]           | et-x/0/19<br>et-x/0/19:<br>[0-1]<br>et-x/0/19:<br>[0-3] | et-x/0/19:<br>[0-2]            | et-x/0/19                      |
|     | et-x/0/20:<br>[0-7]           | et-x/0/20:<br>[0-7]           | et-x/0/20                     | et-x/0/20:<br>[0-8]           | et-x/0/20<br>et-x/0/20:<br>[0-1]<br>et-x/0/20:<br>[0-3] | et-x/0/20:<br>[0-2]            | et-x/0/20                      |
|     | et-x/0/21:<br>[0-7]           | et-x/0/21:<br>[0-7]           | et-x/0/21                     | et-x/0/21:<br>[0-8]           | et-x/0/21<br>et-x/0/21:<br>[0-1]<br>et-x/0/21:<br>[0-3] | et-x/0/21:<br>[0-2]            | et-x/0/21                      |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/22:<br>[0-7]           | et-x/0/22:<br>[0-7]           | et-x/0/22                     | et-x/0/22:<br>[0-8]           | et-x/0/22<br>et-x/0/22:<br>[0-1]<br>et-x/0/22:<br>[0-3] | et-x/0/22:<br>[0-2]            | et-x/0/22                      |
|     | et-x/0/23:<br>[0-7]           | et-x/0/23:<br>[0-7]           | et-x/0/23                     | et-x/0/23:<br>[0-8]           | et-x/0/23<br>et-x/0/23:<br>[0-1]<br>et-x/0/23:<br>[0-3] | et-x/0/23:<br>[0-2]            | et-x/0/23                      |
|     | et-x/0/24:<br>[0-7]           | et-x/0/24:<br>[0-7]           | et-x/0/24                     | et-x/0/24:<br>[0-8]           | et-x/0/24<br>et-x/0/24:<br>[0-1]<br>et-x/0/24:<br>[0-3] | et-x/0/24:<br>[0-2]            | et-x/0/24                      |
|     | et-x/0/25:<br>[0-7]           | et-x/0/25:<br>[0-7]           | et-x/0/25                     | et-x/0/25:<br>[0-8]           | et-x/0/25<br>et-x/0/25:<br>[0-1]<br>et-x/0/25:<br>[0-3] | et-x/0/25:<br>[0-2]            | et-x/0/25                      |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/26:<br>[0-7]           | et-x/0/26:<br>[0-7]           | et-x/0/26                     | et-x/0/26:<br>[0-8]           | et-x/0/26<br>et-x/0/26:<br>[0-1]<br>et-x/0/26:<br>[0-3] | et-x/0/26:<br>[0-2]            | et-x/0/26                      |
|     | et-x/0/27:<br>[0-7]           | et-x/0/27:<br>[0-7]           | et-x/0/27                     | et-x/0/27:<br>[0-8]           | et-x/0/27<br>et-x/0/27:<br>[0-1]<br>et-x/0/27:<br>[0-3] | et-x/0/27:<br>[0-2]            | et-x/0/27                      |
|     | et-x/0/28:<br>[0-7]           | et-x/0/28:<br>[0-7]           | et-x/0/28                     | et-x/0/28:<br>[0-8]           | et-x/0/28<br>et-x/0/28:<br>[0-1]<br>et-x/0/28:<br>[0-3] | et-x/0/28:<br>[0-2]            | et-x/0/28                      |
|     | et-x/0/29:<br>[0-7]           | et-x/0/29:<br>[0-7]           | et-x/0/29                     | et-x/0/29:<br>[0-8]           | et-x/0/29<br>et-x/0/29:<br>[0-1]<br>et-x/0/29:<br>[0-3] | et-x/0/29:<br>[0-2]            | et-x/0/29                      |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/30:<br>[0-7]           | et-x/0/30:<br>[0-7]           | et-x/0/30                     | et-x/0/30:<br>[0-8]           | et-x/0/30<br>et-x/0/30:<br>[0-1]<br>et-x/0/30:<br>[0-3] | et-x/0/30:<br>[0-2]            | et-x/0/30                      |
|     | et-x/0/31:<br>[0-7]           | et-x/0/31:<br>[0-7]           | et-x/0/31                     | et-x/0/31:<br>[0-8]           | et-x/0/31<br>et-x/0/31:<br>[0-1]<br>et-x/0/31:<br>[0-3] | et-x/0/31:<br>[0-2]            | et-x/0/31                      |
|     | et-x/0/32:<br>[0-7]           | et-x/0/32:<br>[0-7]           | et-x/0/32                     | et-x/0/32:<br>[0-8]           | et-x/0/32<br>et-x/0/32:<br>[0-1]<br>et-x/0/32:<br>[0-3] | et-x/0/32:<br>[0-2]            | et-x/0/32                      |
|     | et-x/0/33:<br>[0-7]           | et-x/0/33:<br>[0-7]           | et-x/0/33                     | et-x/0/33:<br>[0-8]           | et-x/0/33<br>et-x/0/33:<br>[0-1]<br>et-x/0/33:<br>[0-3] | et-x/0/33:<br>[0-2]            | et-x/0/33                      |

Table 58: Interface Naming Convention for PTX10K-LC1201 line card (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/0/34:<br>[0-7]           | et-x/0/34:<br>[0-7]           | et-x/0/34                     | et-x/0/34:<br>[0-8]           | et-x/0/34<br>et-x/0/34:<br>[0-1]<br>et-x/0/34:<br>[0-3] | et-x/0/34:<br>[0-2]            | et-x/0/34                      |
|     | et-x/0/35:<br>[0-7]           | et-x/0/35:<br>[0-7]           | et-x/0/35                     | et-x/0/35:<br>[0-8]           | et-x/0/35<br>et-x/0/35:<br>[0-1]<br>et-x/0/35:<br>[0-3] | et-x/0/35:<br>[0-2]            | et-x/0/35                      |

## Port Speed on ACX7100-48L Router Overview

| SUMMARY | IN THIS SECTION                                                                                      |
|---------|------------------------------------------------------------------------------------------------------|
|         | <ul style="list-style-type: none"> <li>Interface Naming Conventions for ACX7100-48L   346</li> </ul> |

For information on ACX7100-48L routers see ACX7100-48L Universal Metro Router Hardware Guide.

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

[Port Speed for ACX7100-48L on page 346](#) summarizes the supported port speeds on a ACX7100 router.

Table 59: Port Speed for ACX7100-48L

| PIC   | Port Number                                                  | Port Speed Supported                                                                                                                                                                                    |
|-------|--------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 | port <b>0 -46</b>                                            | 10-Gigabit Ethernet<br>25-Gigabit Ethernet                                                                                                                                                              |
|       | port 47<br>Port 47 is disabled once PTP function is enabled. | 1-Gigabit Ethernet<br>25-Gigabit Ethernet                                                                                                                                                               |
|       | port <b>48 -53</b>                                           | 100-Mbps and 1-Gigabit Ethernet                                                                                                                                                                         |
| PIC 1 | <b>0 -4</b>                                                  | 400-Gigabit Ethernet<br>1x40-Gigabit Ethernet<br>1x100-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>2x100-Gigabit Ethernet<br>1x100-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>4x25-Gigabit Ethernet |

Use the "[speed](#)" on [page 1137](#) command to set the speed on tri-rate copper SFP port. For information on how to configure the speed at the PIC level, see [Types of Port Speed Configuration on page 295](#).

#### Interface Naming Conventions for ACX7100-48L

Interface Naming Convention for ACX7100-48L lists the interface naming conventions for the ACX7100-48L routers.

**Table 60: Interface Naming Convention for ACX7100-48L**

| PIC   | Interface type                                                                                                                                                 | Interfaces            |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| PIC 0 | 10-Gigabit Ethernet interface (24 SFP5+ ports)                                                                                                                 | et-0/0/0 - et-0/0/47  |
|       | 25-Gigabit Ethernet interface (13 SFP56 ports)                                                                                                                 | et-0/1/24 - et-0/1/47 |
|       | 400-Gigabit/1X40-Gigabit/1X100-Gigabit/4X100-Gigabit/2X100-Gigabit/1X100-Gigabit/8X50-Gigabit/4X10-Gigabit/4X25-Gigabit Ethernet interface (6 QSFP56-DD ports) | et-0/0/48 - et-0/0/53 |

For channelized and non-channelized interface naming formats, see [Channelized and Non-Channelized Interface Naming Formats on page 305](#).

## Interface Naming Conventions for ACX7100-48L

### SUMMARY

Interface Naming Convention for ACX7100-48L lists the interface naming conventions for the ACX7100-48L routers.

**Table 61: Interface Naming Convention for ACX7100-48L**

| PIC   | Interface type                                                                                                                                                 | Interfaces            |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| PIC 0 | 10-Gigabit Ethernet interface (24 SFP5+ ports)                                                                                                                 | et-0/0/0 - et-0/0/47  |
|       | 25-Gigabit Ethernet interface (13 SFP56 ports)                                                                                                                 | et-0/1/24 - et-0/1/47 |
|       | 400-Gigabit/1X40-Gigabit/1X100-Gigabit/4X100-Gigabit/2X100-Gigabit/1X100-Gigabit/8X50-Gigabit/4X10-Gigabit/4X25-Gigabit Ethernet interface (6 QSFP56-DD ports) | et-0/0/48 - et-0/0/53 |

For channelized and non-channelized interface naming formats, see [Channelized and Non-Channelized Interface Naming Formats on page 305](#).

## Port Speed on ACX710 Router Overview

### IN THIS SECTION

- [Interface Naming Conventions | 349](#)

For information on ACX710 routers see *ACX710 Universal Metro Router Hardware Guide*.

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

[Table 62 on page 349](#) summarizes the supported port speeds on a ACX710 router.

Table 62: Port Speed for ACX710

| PIC   | Port Number | Port Speed Supported                                                                                                                                                                 |
|-------|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 | 0 -23       | 10-Gigabit Ethernet<br>1-Gigabit Ethernet                                                                                                                                            |
|       | 0 -15       | 1-Gigabit Ethernet                                                                                                                                                                   |
|       | 16 -23      | 100-Mbps and 1-Gigabit Ethernet                                                                                                                                                      |
| PIC 1 | 0 -4        | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>100-Gigabit Ethernet<br><b>NOTE:</b> By default, all the active ports operate in 100-Gigabit Ethernet mode. |

Use the `"speed"` on page 1137 command to set the speed on tri-rate copper SFP port. For information on how to configure the speed at the PIC level, see Table 33 on page 295. For information on how to configure the speed at the port level, see Table 34 on page 297.

### Interface Naming Conventions

Table 63 on page 349 lists the interface naming conventions for the ACX710 routers.

Table 63: Interface Naming Convention for ACX710

| PIC   | Interface type                                                 | Interfaces           |
|-------|----------------------------------------------------------------|----------------------|
| PIC 0 | 1-Gigabit/10-Gigabit Ethernet interface (24 SFP+ or SFP ports) | xe-0/0/0 – xe-0/0/23 |

**Table 63: Interface Naming Convention for ACX710 (Continued)**

| PIC   | Interface type                                                    | Interfaces            |
|-------|-------------------------------------------------------------------|-----------------------|
|       | 1-Gigabit Ethernet interface (Tri-rate SFP-T optics)              | xe-0/0/0 – xe-0/0/15  |
|       | 100-Mbps and 1-Gigabit Ethernet interface (Tri-rate SFP-T optics) | xe-0/0/16 – xe-0/0/23 |
| PIC 1 | 100-Gigabit Ethernet interface (4 QSFP28 ports)                   | et-0/1/0 – et-0/1/3   |

For channelized and non-channelized interface naming formats, see [Table 37 on page 305](#).

## Port Speed on PTX10001-36MR Router Overview

### IN THIS SECTION

- [Interface Naming Conventions for PTX10001-36MR Router | 356](#)

For information on PTX10001-36MR hardware description, see *PTX10001-36MR Packet Transport Router Hardware Guide*.

[Table 64 on page 350](#) provides the basic details of PTX10001-36MR router.

**Table 64: PTX10001-36MR Router Details and Description**

| Details                               | Description                                       |
|---------------------------------------|---------------------------------------------------|
| PIC Level or Port Level Configuration | Only port level speed configuration is supported. |

Table 64: PTX10001-36MR Router Details and Description (*Continued*)

| Details                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port speed configuration:              | <ul style="list-style-type: none"> <li>• Supports only <b>[edit interfaces interface-name]</b> hierarchy from Junos OS Evolved Release 20.3R1 onwards.</li> <li>• Supports <b>[edit chassis]</b> hierarchy in Junos OS Evolved 20.2 releases.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MTU size                               | <p>PTX10001-36MR router WAN interfaces support maximum MTU of size 16000 bytes for transit traffic. However for the traffic that is destined to host or is originating from host (such as protocol traffic), the maximum MTU limit is 9500 bytes. If any of the host bound packets or host originating packet is above 9500 bytes, then the packet will be dropped. Hence, if you are expecting host packets greater than 9500 bytes, then the WAN interface MTU value must be set as value less than or equal to 9500 bytes.</p>                                                                                                                                                                                                                                                                                                                    |
| Forward Error Correction (FEC) support | <ul style="list-style-type: none"> <li>• By default, KP4 FEC is enabled on 400 and 200-Gigabit Ethernet interfaces.</li> </ul> <p>Since it is mandatory to enable KP4 FEC option for 200 and 400-Gigabit Ethernet interfaces, you cannot disable KP4 FEC explicitly.</p> <ul style="list-style-type: none"> <li>• 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.</li> </ul> <p>For information about FEC support, see <a href="#">"fec (gigether)" on page 713</a>.</p> |

Table 65 on page 352 shows the speed capability of various ports.

**Table 65: 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<br>1/0, 1/1, 1/2, 1/3, 1/8, 1/9, 1/10, 1/11<br>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<br>1/4, 1/5, 1/6, 1/7<br>2/4, 2/5, 2/6, 2/7                                                                   | 100-Gbps         |

[Table 66 on page 352](#) summarizes the channelization and FEC support on 400- and 200-Gigabit Ethernet capable ports.

**Table 66: Channelization and FEC support on 400- and 200-Gigabit Ethernet capable ports**

| Speed Supported                    | You can channelize to:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 400-Gigabit Ethernet capable ports | <ul style="list-style-type: none"> <li>• One or two 200-Gigabit Ethernet interfaces (1x200G or 2x200G)</li> <li>• Four 100-Gigabit Ethernet interfaces or two 100-Gigabit Ethernet interfaces (4x100G or 2x100G)</li> <li>• Two 50-Gigabit Ethernet interfaces or eight 50-Gigabit Ethernet interfaces (2x50G or 8x50G)</li> <li>• Four 25-Gigabit Ethernet interfaces or eight 25-Gigabit Ethernet interfaces (4x25G or 8x25G)</li> <li>• One 40-Gigabit Ethernet interface (1x40G)</li> <li>• Four 10-Gigabit Ethernet interfaces (4x10G)</li> </ul> |

**Table 66: Channelization and FEC support on 400- and 200-Gigabit Ethernet capable ports (Continued)**

| Speed Supported                    | You can channelize to:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 100-Gigabit Ethernet capable ports | <p>You can configure all four 100-Gigabit Ethernet capable ports (ports 4, 5, 6, and 7) to operate at 100-Gbps speeds.</p> <p>Note that, you can configure port 4 and port 6 to operate at 1x100-Gbps, 4x10-Gbps, 4x25-Gbps, or 1x40-Gbps speeds with the following conditions:</p> <ul style="list-style-type: none"> <li>• When port 4 is configured in 4x10-Gbps, 4x25-Gbps, or 1x40-Gbps, then port 5 must be configured as 'unused'.</li> <li>• When port 6 is configured in 4x10-Gbps, 4x25-Gbps, or 1x40-Gbps, then port 7 must be configured as 'unused'.</li> </ul> |

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.

### Configure speed at Port Level

To configure the PTX10001-36MR router at port level, follow the configuration steps in [Table 34 on page 297](#) in 20.2 releases. See "[speed](#)" on [page 1134](#) for more details.

From Junos OS Evolved Release 20.3R1 you must use, **set interfaces et-*<fpc>*/*<pic>*/*<port>* speed <10g | 25g | 50g | 40g | 100g | 200g | 400g>**, or **set interfaces et-*<fpc>*/*<pic>*/*<port>* speed <10g | 25g | 50g | 40g | 100g | 200g | 400g> number-of-sub-ports (1|2|4|8)** commands to configure the speed on the ports.

You can configure port profiles in the command line interface without the physical presence of an FPC. If an invalid port profile configuration is detected while booting a FPC, an alarm is generated. Also, the default port profile is selected for that PIC.

If the port profile configuration is changed while the FPC is up and running, and the new configuration is invalid, an alarm is generated. The existing port profile configured continues to be used for that PIC.

You can configure any supported speed on 400-Gbps capable ports. Configuring speed on one of the 400-Gbps capable port will not disrupt the traffic on any other ports.

But, for 100-Gbps capable ports only speeds of 100-Gbps, 25-Gbps, 40-Gbps, and 10-Gbps are valid.

You must explicitly configure the ports that must be powered off, using the following command:

**set chassis fpc *fpc-slot* pic *pic-slot* port *port-num* unused.**

From Junos OS Evolved Release 20.3R1 you must use, the following command:

**set interfaces *et-*<fpc>/<pic>/<port>** unused**

If you configure "Unused" on page 1220 command for a physical port in a PIC, then no channelized and non-channelized interfaces will be created. See "Unused" on page 1220 for more information.

The Table 69 on page 356 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 67 on page 354 provides you the ports that you must power-off while configuring different speeds.

**Table 67: 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<br>0/1/4 and 0/1/6<br>0/2/4 and 0/2/6                                                                         | 0/0/5 and 0/0/7<br>0/1/5 and 0/1/7<br>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.

To specify which of these channelization modes the port should operate in, execute the following command:

**set chassis fpc *fpc-slot* pic *pic-slot* port *port-num* *number-of-sub-ports* (1|2|4|8)**

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](#)" on page 1042 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 68 on page 355](#):

**Table 68: Number of sub-ports supported for a particular speed**

| Port speed | Valid values for <b>number-of-sub-ports</b> |
|------------|---------------------------------------------|
| 200-Gbps   | 1, 2                                        |
| 100-Gbps   | 1, 2, 4                                     |
| 40-Gbps    | 1                                           |
| 50-Gbps    | 2, 8                                        |
| 25-Gbps    | 4, 8                                        |
| 10-Gbps    | 4                                           |

Follow these guidelines when you configure the speed of a port:

- The default number of channels is considered as 1 for all speeds if number of sub-ports is not configured explicitly.
- 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. For example, if a port is channelized for eight channels and the optics inserted is for four channels, then the links will be down.
- If the **number-of-subports** configuration is incorrect, then the error will be logged in syslog and an alarm is raised.

- 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 mode.
- To avoid oversubscription of a PIC, the sum of total speed configured on ports 4 to 11 within the PIC must be less than 1.6 Tbps.

Use the **show chassis pic fpc-slot 0 pic-slot 0** command in 20.2 release at the **[edit chassis]** hierarchy to know the speed that each port is configured for in the PIC. From Junos OS Evolved Release 20.3R1 you must use the operational command "[show interfaces port-profile](#)" on page 1695.

### Interface Naming Conventions for PTX10001-36MR Router

[Table 69 on page 356](#) lists the interface naming conventions for the PTX10001-36MR router.

**Table 69: Interface Naming Convention for PTX10001-36MR Router**

| PIC | 4x100-Gigabit Ethernet Interface | 4x10-Gigabit Ethernet Interface<br>4x25-Gigabit Ethernet Interface | 2x100-Gigabit Ethernet Interface<br>2x50-Gigabit Ethernet Interface<br>2x200-Gigabit Ethernet Interface | 8x25-Gigabit Ethernet Interface<br>8x50-Gigabit Ethernet Interface | 1x400-Gigabit Ethernet Interface<br>1x200-Gigabit Ethernet Interface | 1x100-Gigabit Ethernet Interface | 1x40-Gigabit Ethernet Interface |
|-----|----------------------------------|--------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------|----------------------------------------------------------------------|----------------------------------|---------------------------------|
| 0   | et-0/0/0:<br>[0-3]               | et-0/0/0:<br>[0-3]                                                 | et-0/0/0:<br>[0-1]                                                                                      | et-0/0/0:<br>[0-7]                                                 | et-0/0/0                                                             | et-0/0/0                         | et-0/0/0                        |
|     | et-0/0/1:<br>[0-3]               | et-0/0/1:<br>[0-3]                                                 | et-0/0/1:<br>[0-1]                                                                                      | et-0/0/1:<br>[0-7]                                                 | et-0/0/1                                                             | et-0/0/1                         | et-0/0/1                        |
|     | et-0/0/2:<br>[0-3]               | et-0/0/2:<br>[0-3]                                                 | et-0/0/2:<br>[0-1]                                                                                      | et-0/0/2:<br>[0-7]                                                 | et-0/0/2                                                             | et-0/0/2                         | et-0/0/2                        |

Table 69: Interface Naming Convention for PTX10001-36MR Router (Continued)

|                        |                                  |                                 |                                  |                                 |                                  |                                  |                                 |
|------------------------|----------------------------------|---------------------------------|----------------------------------|---------------------------------|----------------------------------|----------------------------------|---------------------------------|
| PIC                    | 4x100-Gigabit Ethernet Interface | 4x10-Gigabit Ethernet Interface | 2x100-Gigabit Ethernet Interface | 8x25-Gigabit Ethernet Interface | 1x400-Gigabit Ethernet Interface | 1x100-Gigabit Ethernet Interface | 1x40-Gigabit Ethernet Interface |
|                        |                                  | 4x25-Gigabit Ethernet Interface | 2x50-Gigabit Ethernet Interface  | 8x50-Gigabit Ethernet Interface | 1x200-Gigabit Ethernet Interface |                                  |                                 |
|                        |                                  |                                 | 2x200-Gigabit Ethernet Interface |                                 |                                  |                                  |                                 |
|                        | <b>et-0/0/3: [0-3]</b>           | <b>et-0/0/3: [0-3]</b>          | <b>et-0/0/3: [0-1]</b>           | <b>et-0/0/3: [0-7]</b>          | <b>et-0/0/3</b>                  | <b>et-0/0/3</b>                  | <b>et-0/0/3</b>                 |
|                        | Unsupported                      | <b>et-0/0/4: [0-3]</b>          | Unsupported                      | Unsupported                     | Unsupported                      | <b>et-0/0/4</b>                  | <b>et-0/0/4</b>                 |
|                        | Unsupported                      | Unused                          | Unsupported                      | Unsupported                     | Unsupported                      | <b>et-0/0/5</b>                  | Unused                          |
|                        | Unsupported                      | <b>et-0/0/6: [0-3]</b>          | Unsupported                      | Unsupported                     | Unsupported                      | <b>et-0/0/6</b>                  | <b>et-0/0/6</b>                 |
|                        | Unsupported                      | Unused                          | Unsupported                      | Unsupported                     | Unsupported                      | <b>et-0/0/7</b>                  | Unused                          |
| <b>et-0/0/8: [0-3]</b> | <b>et-0/0/8: [0-3]</b>           | <b>et-0/0/8: [0-1]</b>          | <b>et-0/0/8: [0-7]</b>           | <b>et-0/0/8</b>                 | <b>et-0/0/8</b>                  | <b>et-0/0/8</b>                  |                                 |
| <b>et-0/0/9: [0-3]</b> | <b>et-0/0/9: [0-3]</b>           | <b>et-0/0/9: [0-1]</b>          | <b>et-0/0/9: [0-7]</b>           | <b>et-0/0/9</b>                 | <b>et-0/0/9</b>                  | <b>et-0/0/9</b>                  |                                 |

Table 69: Interface Naming Convention for PTX10001-36MR Router (Continued)

| PIC | 4x100-Gigabit Ethernet Interface | 4x10-Gigabit Ethernet Interface | 2x100-Gigabit Ethernet Interface | 8x25-Gigabit Ethernet Interface | 1x400-Gigabit Ethernet Interface | 1x100-Gigabit Ethernet Interface | 1x40-Gigabit Ethernet Interface |
|-----|----------------------------------|---------------------------------|----------------------------------|---------------------------------|----------------------------------|----------------------------------|---------------------------------|
|     |                                  | 4x25-Gigabit Ethernet Interface | 2x50-Gigabit Ethernet Interface  | 8x50-Gigabit Ethernet Interface | 1x200-Gigabit Ethernet Interface |                                  |                                 |
|     |                                  |                                 | 2x200-Gigabit Ethernet Interface |                                 |                                  |                                  |                                 |
|     | et-0/0/10:<br>[0-3]              | et-0/0/10:<br>[0-3]             | et-0/0/10:<br>[0-1]              | et-0/0/10:<br>[0-7]             | et-0/0/10                        | et-0/0/10                        | et-0/0/10                       |
|     | et-0/0/11:<br>[0-3]              | et-0/0/11:<br>[0-3]             | et-0/0/11:<br>[0-1]              | et-0/0/11:<br>[0-7]             | et-0/0/11                        | et-0/0/11                        | et-0/0/11                       |
| 1   | et-0/1/0:<br>[0-3]               | et-0/1/0:<br>[0-3]              | et-0/1/0:<br>[0-1]               | et-0/1/0:<br>[0-7]              | et-0/1/0                         | et-0/1/0                         | et-0/1/0                        |
|     | et-0/1/1:<br>[0-3]               | et-0/1/1:<br>[0-3]              | et-0/1/1:<br>[0-1]               | et-0/1/1:<br>[0-7]              | et-0/1/1                         | et-0/1/1                         | et-0/1/1                        |
|     | et-0/1/2:<br>[0-3]               | et-0/1/2:<br>[0-3]              | et-0/1/2:<br>[0-1]               | et-0/1/2:<br>[0-7]              | et-0/1/2                         | et-0/1/2                         | et-0/1/2                        |
|     | et-0/1/3:<br>[0-3]               | et-0/1/3:<br>[0-3]              | et-0/1/3:<br>[0-1]               | et-0/1/3:<br>[0-7]              | et-0/1/3                         | et-0/1/3                         | et-0/1/3                        |
|     | Unsuppor<br>ted                  | et-0/1/4:<br>[0-3]              | Unsuppor<br>ted                  | Unsuppor<br>ted                 | Unsuppor<br>ted                  | et-0/1/4                         | et-0/1/4                        |

Table 69: Interface Naming Convention for PTX10001-36MR Router (Continued)

| PIC | 4x100-Gigabit Ethernet Interface | 4x10-Gigabit Ethernet Interface<br>4x25-Gigabit Ethernet Interface | 2x100-Gigabit Ethernet Interface<br>2x50-Gigabit Ethernet Interface<br>2x200-Gigabit Ethernet Interface | 8x25-Gigabit Ethernet Interface<br>8x50-Gigabit Ethernet Interface | 1x400-Gigabit Ethernet Interface<br>1x200-Gigabit Ethernet Interface | 1x100-Gigabit Ethernet Interface | 1x40-Gigabit Ethernet Interface |
|-----|----------------------------------|--------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------|----------------------------------------------------------------------|----------------------------------|---------------------------------|
|     | Unsupported                      | Unused                                                             | Unsupported                                                                                             | Unsupported                                                        | Unsupported                                                          | <b>et-0/1/5</b>                  | Unused                          |
|     | Unsupported                      | <b>et-0/1/6:<br/>[0-3]</b>                                         | Unsupported                                                                                             | Unsupported                                                        | Unsupported                                                          | <b>et-0/1/6</b>                  | <b>et-0/1/6</b>                 |
|     | Unsupported                      | Unused                                                             | Unsupported                                                                                             | Unsupported                                                        | Unsupported                                                          | <b>et-0/1/7</b>                  | Unused                          |
|     | <b>et-0/1/8:<br/>[0-3]</b>       | <b>et-0/1/8:<br/>[0-3]</b>                                         | <b>et-0/1/8:<br/>[0-1]</b>                                                                              | <b>et-0/1/8:<br/>[0-7]</b>                                         | <b>et-0/1/8</b>                                                      | <b>et-0/1/8</b>                  | <b>et-0/1/8</b>                 |
|     | <b>et-0/1/9:<br/>[0-3]</b>       | <b>et-0/1/9:<br/>[0-3]</b>                                         | <b>et-0/1/9:<br/>[0-1]</b>                                                                              | <b>et-0/1/9:<br/>[0-7]</b>                                         | <b>et-0/1/9</b>                                                      | <b>et-0/1/9</b>                  | <b>et-0/1/9</b>                 |
|     | <b>et-0/1/10:<br/>:[0-3]</b>     | <b>et-0/1/10:<br/>:[0-3]</b>                                       | <b>et-0/1/10:<br/>:[0-1]</b>                                                                            | <b>et-0/1/10:<br/>:[0-7]</b>                                       | <b>et-0/1/10</b>                                                     | <b>et-0/1/10</b>                 | <b>et-0/1/10</b>                |
|     | <b>et-0/1/11:<br/>:[0-3]</b>     | <b>et-0/1/11:<br/>:[0-3]</b>                                       | <b>et-0/1/11:<br/>:[0-1]</b>                                                                            | <b>et-0/1/11:<br/>:[0-7]</b>                                       | <b>et-0/1/11</b>                                                     | <b>et-0/1/11</b>                 | <b>et-0/1/11</b>                |

Table 69: Interface Naming Convention for PTX10001-36MR Router (Continued)

| PIC | 4x100-Gigabit Ethernet Interface | 4x10-Gigabit Ethernet Interface<br><br>4x25-Gigabit Ethernet Interface | 2x100-Gigabit Ethernet Interface<br><br>2x50-Gigabit Ethernet Interface<br><br>2x200-Gigabit Ethernet Interface | 8x25-Gigabit Ethernet Interface<br><br>8x50-Gigabit Ethernet Interface | 1x400-Gigabit Ethernet Interface<br><br>1x200-Gigabit Ethernet Interface | 1x100-Gigabit Ethernet Interface | 1x40-Gigabit Ethernet Interface |
|-----|----------------------------------|------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|--------------------------------------------------------------------------|----------------------------------|---------------------------------|
| 2   | et-0/2/0:<br>[0-3]               | et-0/2/0:<br>[0-3]                                                     | et-0/2/0:<br>[0-1]                                                                                              | et-0/2/0:<br>[0-7]                                                     | et-0/2/0                                                                 | et-0/2/0                         | et-0/2/0                        |
|     | et-0/2/1:<br>[0-3]               | et-0/2/1:<br>[0-3]                                                     | et-0/2/1:<br>[0-1]                                                                                              | et-0/2/1:<br>[0-7]                                                     | et-0/2/1                                                                 | et-0/2/1                         | et-0/2/1                        |
|     | et-0/2/2:<br>[0-3]               | et-0/2/2:<br>[0-3]                                                     | et-0/2/2:<br>[0-1]                                                                                              | et-0/2/2:<br>[0-7]                                                     | et-0/2/2                                                                 | et-0/2/2                         | et-0/2/2                        |
|     | et-0/2/3:<br>[0-3]               | et-0/2/3:<br>[0-3]                                                     | et-0/2/3:<br>[0-1]                                                                                              | et-0/2/3:<br>[0-7]                                                     | et-0/2/3                                                                 | et-0/2/3                         | et-0/2/3                        |
|     | Unsupported                      | et-0/2/4:<br>[0-3]                                                     | Unsupported                                                                                                     | Unsupported                                                            | Unsupported                                                              | et-0/2/4                         | et-0/2/4                        |
|     | Unsupported                      | Unused                                                                 | Unsupported                                                                                                     | Unsupported                                                            | Unsupported                                                              | et-0/2/5                         | Unused                          |
|     | Unsupported                      | et-0/2/6:<br>[0-3]                                                     | Unsupported                                                                                                     | Unsupported                                                            | Unsupported                                                              | et-0/2/6                         | et-0/2/6                        |

Table 69: Interface Naming Convention for PTX10001-36MR Router (Continued)

| PIC | 4x100-Gigabit Ethernet Interface | 4x10-Gigabit Ethernet Interface | 2x100-Gigabit Ethernet Interface | 8x25-Gigabit Ethernet Interface | 1x400-Gigabit Ethernet Interface | 1x100-Gigabit Ethernet Interface | 1x40-Gigabit Ethernet Interface |
|-----|----------------------------------|---------------------------------|----------------------------------|---------------------------------|----------------------------------|----------------------------------|---------------------------------|
|     |                                  | 4x25-Gigabit Ethernet Interface | 2x50-Gigabit Ethernet Interface  | 8x50-Gigabit Ethernet Interface | 1x200-Gigabit Ethernet Interface |                                  |                                 |
|     |                                  |                                 | 2x200-Gigabit Ethernet Interface |                                 |                                  |                                  |                                 |
|     | Unsupported                      | Unused                          | Unsupported                      | Unsupported                     | Unsupported                      | et-0/2/7                         | Unused                          |
|     | et-0/2/8:[0-3]                   | et-0/2/8:[0-3]                  | et-0/2/8:[0-1]                   | et-0/2/8:[0-7]                  | et-0/2/8                         | et-0/2/8                         | et-0/2/8                        |
|     | et-0/2/9:[0-3]                   | et-0/2/9:[0-3]                  | et-0/2/9:[0-1]                   | et-0/2/9:[0-7]                  | et-0/2/9                         | et-0/2/9                         | et-0/2/9                        |
|     | 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                        |                                 |

Refer to ["Interface Naming Conventions"](#) on page 304 for channelized and non-channelized interface naming formats.

### Port Speed on PTX10K-LC1202-36MR Overview

For information on the line card, see [PTX10008 Line Card Components and Descriptions](#).

For information about platform support, see [Hardware Compatibility Tool \(HCT\)](#).

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.

[Table 70 on page 362](#) summarizes the supported port speeds on PTX10K-LC1202-36MR for PTX10008 Routers.

**Table 70: 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<br>4x100-Gigabit Ethernet<br>2x100-Gigabit Ethernet<br>40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>8x25-Gigabit Ethernet<br>10-Gigabit Ethernet<br><b>Default Speed:</b> 400-Gigabit Ethernet |
|                     | 100G capable ports - 0 - 35            | 100-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>40-Gigabit Ethernet<br>10-Gigabit Ethernet<br><b>Default Speed:</b> 100-Gigabit Ethernet                                                     |

[Table 35 on page 298](#) and [Table 36 on page 301](#) describe 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" on page 1137](#) 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" on page 1042](#).

When the `number-of-sub-ports` are not specified, the number of channels are created as per [Table 71 on page 363](#):

**Table 71: Number of sub-ports supported for a particular speed**

| Port speed | Valid values for <code>number-of-sub-ports</code> |
|------------|---------------------------------------------------|
| 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
```

See [Table 36 on page 301](#) and ["Unused" on page 1220](#) for more details.

- 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 \(gether\)" on page 713](#).

You can configure every interface to loopback mode, see ["loopback" on page 984](#).

[Table 72 on page 364](#) lists the naming conventions used for interfaces on PTX10K-LC1202-36MR (for 100-Gbps, 2x100-Gbps, 400-Gbps, and 4x100-Gbps speeds) for PTX10008 routers. PTX10008 routers support eight PTX10K-LC1202-36MR line cards.

**Table 72: 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 | <b>et-x/0/0</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|       | <b>et-x/0/1</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|       | <b>et-x/0/2</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|       | <b>et-x/0/3</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|       | <b>et-x/0/4</b>                | <b>et-x/0/4:[0-1]</b>            | <b>et-x/0/4</b>                | <b>et-x/0/4:[0-3]</b>            |
|       | <b>et-x/0/5</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|       | <b>et-x/0/6</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|       | <b>et-x/0/7</b>                | Not Supported                    | Not Supported                  | Not Supported                    |

**Table 72: 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 |
|-----|--------------------------------|----------------------------------|--------------------------------|----------------------------------|
|     | <b>et-x/0/8</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/9</b>                | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/10</b>               | <b>et-x/0/10:[0-1]</b>           | <b>et-x/0/10</b>               | <b>et-x/0/10:[0-3]</b>           |
|     | <b>et-x/0/11</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/12</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/13</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/14</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/15</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/16</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/17</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/18</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/19</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/20</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/21</b>               | Not Supported                    | Not Supported                  | Not Supported                    |

**Table 72: 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 |
|-----|--------------------------------|----------------------------------|--------------------------------|----------------------------------|
|     | <b>et-x/0/22</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/23</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/24</b>               | <b>et-x/0/24:[0-1]</b>           | <b>et-x/0/24</b>               | <b>et-x/0/24:[0-3]</b>           |
|     | <b>et-x/0/25</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/26</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/27</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/28</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/29</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/30</b>               | <b>et-x/0/30:[0-1]</b>           | <b>et-x/0/30</b>               | <b>et-x/0/30:[0-3]</b>           |
|     | <b>et-x/0/31</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/32</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/33</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/34</b>               | Not Supported                    | Not Supported                  | Not Supported                    |
|     | <b>et-x/0/35</b>               | Not Supported                    | Not Supported                  | Not Supported                    |

Table 73 on page 367 lists the naming conventions used for interfaces on PTX10K-LC1202-36MR (for 10-Gbps, 4x10-Gbps, 40-Gbps, 4x25-Gbps, and 8x25-Gbps speeds) for PTX10008 routers. PTX10008 routers support eight PTX10K-LC1202-36MR line cards.

**Table 73: 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 | <b>et-x/0/0</b>               | <b>et-x/0/0:[0-3]</b>           | <b>et-x/0/0</b>               | <b>et-x/0/0:[0-3]</b>           | Not Supported                   |
|       | Not Supported                 | Not Supported                   | Not Supported                 | Not Supported                   | Not Supported                   |
|       | <b>et-x/0/2</b>               | <b>et-x/0/2:[0-3]</b>           | <b>et-x/0/2</b>               | <b>et-x/0/2:[0-3]</b>           | Not Supported                   |
|       | Not Supported                 | Not Supported                   | Not Supported                 | Not Supported                   | Not Supported                   |
|       | <b>et-x/0/4</b>               | <b>et-x/0/4:[0-3]</b>           | <b>et-x/0/4</b>               | <b>et-x/0/4:[0-3]</b>           | <b>et-x/0/4:[0-7]</b>           |
|       | <b>et-x/0/5</b>               | <b>et-x/0/5:[0-3]</b>           | <b>et-x/0/5</b>               | <b>et-x/0/5:[0-3]</b>           | Not Supported                   |
|       | <b>et-x/0/6</b>               | <b>et-x/0/6:[0-3]</b>           | <b>et-x/0/6</b>               | <b>et-x/0/6:[0-3]</b>           | Not Supported                   |
|       | <b>et-x/0/7</b>               | <b>et-x/0/7:[0-3]</b>           | <b>et-x/0/7</b>               | <b>et-x/0/7:[0-3]</b>           | Not Supported                   |
|       | <b>et-x/0/8</b>               | <b>et-x/0/8:[0-3]</b>           | <b>et-x/0/8</b>               | <b>et-x/0/8:[0-3]</b>           | Not Supported                   |
|       | <b>et-x/0/9</b>               | <b>et-x/0/9:[0-3]</b>           | <b>et-x/0/9</b>               | <b>et-x/0/9:[0-3]</b>           | Not Supported                   |
|       | <b>et-x/0/10</b>              | <b>et-x/0/10:[0-3]</b>          | <b>et-x/0/10</b>              | <b>et-x/0/10:[0-3]</b>          | <b>et-x/0/10:<br/>[0-7]</b>     |
|       | <b>et-x/0/11</b>              | <b>et-x/0/11:[0-3]</b>          | <b>et-x/0/11</b>              | <b>et-x/0/11:[0-3]</b>          | Not Supported                   |

**Table 73: 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/0/12                     | et-x/0/12:[0-3]                 | et-x/0/12                     | et-x/0/12:[0-3]                 | Not Supported                   |
|     | et-x/0/13                     | et-x/0/13:[0-3]                 | et-x/0/13                     | et-x/0/13:[0-3]                 | Not Supported                   |
|     | et-x/0/14                     | et-x/0/14:[0-3]                 | et-x/0/14                     | et-x/0/14:[0-3]                 | Not Supported                   |
|     | et-x/0/15                     | et-x/0/15:[0-3]                 | et-x/0/15                     | et-x/0/15:[0-3]                 | Not Supported                   |
|     | et-x/0/16                     | et-x/0/16:[0-3]                 | et-x/0/16                     | et-x/0/16:[0-3]                 | Not Supported                   |
|     | et-x/0/17                     | et-x/0/17:[0-3]                 | et-x/0/17                     | et-x/0/17:[0-3]                 | Not Supported                   |
|     | et-x/0/18                     | et-x/0/18:[0-3]                 | et-x/0/18                     | et-x/0/18:[0-3]                 | Not Supported                   |
|     | Not Supported                 | Not Supported                   | Not Supported                 | Not Supported                   | Not Supported                   |
|     | et-x/0/20                     | et-x/0/20:[0-3]                 | et-x/0/20                     | et-x/0/20:[0-3]                 | Not Supported                   |
|     | Not Supported                 | Not Supported                   | Not Supported                 | Not Supported                   | Not Supported                   |
|     | et-x/0/22                     | et-x/0/22:[0-3]                 | et-x/0/22                     | et-x/0/22:[0-3]                 | Not Supported                   |
|     | et-x/0/23                     | et-x/0/23:[0-3]                 | et-x/0/23                     | et-x/0/23:[0-3]                 | Not Supported                   |
|     | et-x/0/24                     | et-x/0/24:[0-3]                 | et-x/0/24                     | et-x/0/24:[0-3]                 | et-x/0/24:<br>[0-7]             |

**Table 73: 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/0/25                     | et-x/0/25:[0-3]                 | et-x/0/25                     | et-x/0/25:[0-3]                 | Not Supported                   |
|     | et-x/0/26                     | et-x/0/26:[0-3]                 | et-x/0/26                     | et-x/0/26:[0-3]                 | Not Supported                   |
|     | et-x/0/27                     | et-x/0/27:[0-3]                 | et-x/0/27                     | et-x/0/27:[0-3]                 | Not Supported                   |
|     | et-x/0/28                     | et-x/0/28:[0-3]                 | et-x/0/28                     | et-x/0/28:[0-3]                 | Not Supported                   |
|     | et-x/0/29                     | et-x/0/29:[0-3]                 | et-x/0/29                     | et-x/0/29:[0-3]                 | Not Supported                   |
|     | et-x/0/30                     | et-x/0/30:[0-3]                 | et-x/0/30                     | et-x/0/30:[0-3]                 | et-x/0/30:<br>[0-7]             |
|     | et-x/0/31                     | et-x/0/31:[0-3]                 | et-x/0/31                     | et-x/0/31:[0-3]                 | Not Supported                   |
|     | et-x/0/32                     | et-x/0/32:[0-3]                 | et-x/0/32                     | et-x/0/32:[0-3]                 | Not Supported                   |
|     | et-x/0/33                     | et-x/0/33:[0-3]                 | et-x/0/33                     | et-x/0/33:[0-3]                 | Not Supported                   |
|     | et-x/0/34                     | et-x/0/34:[0-3]                 | et-x/0/34                     | et-x/0/34:[0-3]                 | Not Supported                   |
|     | et-x/0/35                     | et-x/0/35:[0-3]                 | et-x/0/35                     | et-x/0/35:[0-3]                 | Not Supported                   |

See ["Interface Naming Conventions"](#) on page 304 for channelized and non-channelized interface naming formats.

## Introduction to Port Speed

### IN THIS SECTION

- [Port Speed for Routing Devices | 370](#)
- [Guidelines for Configuring Port Speed on Routing Devices | 374](#)
- [MX10003 MPC Port Speed Overview | 376](#)
- [MX204 Router Port Speed Overview | 382](#)
- [PTX10003 Router Port Speed Overview | 390](#)
- [PTX10K-LC1201 Port Speed Overview | 392](#)

Use this topic to understand more about port speed in a network device or in a network component such as a line card.

### Port Speed for Routing Devices

### IN THIS SECTION

- [Port Speed on MPC7E-MRATE | 370](#)
- [Port Speed on MIC-MRATE | 371](#)
- [Port Speed on PTX10K-LC2101 | 372](#)
- [Port Speed on MIC-MACSEC-20GE | 373](#)

The total port speed of the MIC/MPC cannot exceed the forwarding capacity of the Packet Forwarding Engine.

#### *Port Speed on MPC7E-MRATE*

MPC7E (MPC7E-MRATE) is a fixed-configuration MPC and is supported on MX240, MX480, MX960, MX2010, and MX2020 routers. MPC7E-MRATE contains two built-in PICs, PIC 0 and PIC 1. Each PIC has six physical ports that support quad small form-factor pluggable plus (QSFP+) transceivers. The default port speed is 10 Gbps for all ports. Each of the six ports of PIC 0 and PIC 1 supports speeds of 10 Gbps and 40 Gbps. However, only ports 2 and 5 on both the PICs support 100 Gbps speed.

MPC7E-MRATE has an aggregate forwarding capacity of 480 Gbps and a forwarding capacity of 240 Gbps on each Packet Forwarding Engine. Oversubscription of Packet Forwarding Engine capacity is not supported. The demand on each Packet Forwarding Engine must be less than or equal to its forwarding capacity. For instance, for MPC7E-MRATE, the demand on each Packet Forwarding Engine must be less than or equal to 240 Gbps.

For information about the naming conventions for interfaces on MPC7E-MRATE MPC, see *Interface Naming Conventions for MPC7E-MRATE*.

Starting with Junos OS Release 20.4R1, you can now configure 1-Gbps speed on 10-Gigabit Ethernet ports of the MPC7E-10G.

Each of the forty 10-Gigabit Ethernet ports can be configured to operate as 1-Gigabit Ethernet port. To configure the operating speed of the 10-Gbps port to 1-Gbps, use the **speed** statement at the **edit interfaces *interfacename* gigheter-options** hierarchy level. After you commit the configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed without any MPC, PIC, or interface bounce. To view the speed configured on the interface, use the **show interfaces extensive** command. In the output, the **Speed Configuration** field displays the current operating speed of the interface. If the interface is configured with 1-Gbps speed, then the **Speed Configuration** field displays 1G; if the interface is configured with 10-Gbps speed, **Speed Configuration** displays AUTO. For more information, see *speed*. Use the `show interfaces extensive` command to view the speed of the interface.

### *Port Speed on MIC-MRATE*

MPC8E (MX2K-MPC8E) and MPC9E (MX2K-MPC9E) support two separate slots for MICs as field replaceable units (FRUs). Each of the MIC slots supports only one MIC—MIC-MRATE. MIC-MRATE consists of 12 physical ports that support QSFP+ transceivers and multiple port speeds of 100 Gbps, 40 Gbps, and 10 Gbps. You can configure a port to operate in a specific speed based on your requirement. The default port speed is 10 Gbps for all ports. MIC-MRATE also supports breakout transceivers, which you can use to split a 40-Gigabit Ethernet port into four 10-Gigabit Ethernet ports. MIC-MRATE ports can be split into a maximum of 48 10-Gigabit Ethernet interfaces.

MPC8E has an aggregate forwarding capacity of 960 Gbps and a forwarding capacity of 240 Gbps on each Packet Forwarding Engine. MPC9E has an aggregate forwarding capacity of 1600 Gbps and a forwarding capacity of 400 Gbps on each Packet Forwarding Engine. Oversubscription of Packet Forwarding Engine capacity is not supported. The demand on each Packet Forwarding Engine must be less than or equal to its forwarding capacity. For instance, for MPC8E, the demand on each Packet Forwarding Engine must be less than or equal to 240 Gbps and for MPC9E, the demand per Packet Forwarding Engine must be less than or equal to 400 Gbps.

On MPC8E with MIC-MRATE, you can configure four 100-Gigabit Ethernet interfaces and two 40-Gigabit Ethernet interfaces per MIC. All other interfaces are automatically disabled. On MPC9E with MIC-MRATE, you can configure eight ports as 100-Gigabit Ethernet interfaces and the other ports can be configured only as 40-Gigabit Ethernet Interfaces or 10-Gigabit Ethernet interfaces.

For information about the naming conventions for interfaces on MPC8E and MPC9E, see *Interface Naming Conventions for MIC-MRATE*.

### ***Port Speed on PTX10K-LC2101***

PTX10K-LC2101 is a fixed-configuration MPC and is supported on MX10008 routers. PTX10K-LC2101 contains six built-in PICs, PIC 0 to PIC 5. Each PIC has four physical ports that support quad small form-factor pluggable plus (QSFP+) transceivers. The default port speed is 10 Gbps for all ports. Each of the four ports of PIC 0 to PIC 5 supports speeds of 10 Gbps (using breakout cables), 40 Gbps, and 100 Gbps.

MX10008 routers support eight PTX10K-LC2101 MPCs. By default, each PTX10K-LC2101 MPC provides a maximum bandwidth of 1.44 Tbps. PTX10K-LC2101 has six Packet Forwarding Engines, each providing a maximum bandwidth of up to 240 Gbps, which cannot be oversubscribed. You can configure PTX10K-LC2101 to provide an increased bandwidth of 2.4 Tbps. The demand on each Packet Forwarding Engine must be less than or equal to its forwarding capacity. For instance, by default, for PTX10K-LC2101, the demand on each Packet Forwarding Engine must be less than or equal to 240 Gbps. However, if you have configured PTX10K-LC2101 to provide an increased bandwidth of 2.4 Tbps, the demand on each Packet Forwarding Engine must be less than or equal to 400 Gbps.

For information about the naming conventions for interfaces on PTX10K-LC2101 MPC, see *Interface Naming Conventions for PTX10K-LC2101*.

Starting with Junos OS Release 19.4R1, you can now configure 1-Gbps speed on 10-Gigabit Ethernet ports of the JNP10K-LC2101 MPC.

Each of the 40-Gigabit Ethernet port can be split to four 10-Gigabit Ethernet ports that can be configured to operate as 1-Gigabit Ethernet port. You must use 4x10GE LR breakout optics (QSFP-4X10GE-LR) at the MX10008 or MX10016 end and 1-Gigabit Ethernet EX optics at the remote end. It is only optional to use Juniper optics (SFP-GE40KM) at the remote end, as any vendor's EX (not SX or LX) optics can be used. Refer to the [Hardware Compatibility Tool](#) for the list of pluggable transceivers supported on the MX10008 router.

**NOTE:** Any interface operating at 10-Gbps speed can be independently converted to 1-Gbps speed. For example, in multi-rate connections through split cables, when one of the ports operates at 1-Gbps speed, the other three ports can be configured either with 1-Gbps speed or 10-Gbps speed.

To configure the operating speed of the 10-Gbps port to 1-Gbps, use the **speed** statement at the **edit interfaces *interfacename* together-options** hierarchy level. After you commit the configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed without any MPC, PIC, or interface bounce. To view the speed configured on the interface, use the **show interfaces extensive** command. In the output, the **Speed Configuration** field displays the current operating speed of the interface. If the

interface is configured with 1-Gbps speed, then the **Speed Configuration** field displays 1G; if the interface is configured with 10-Gbps speed, **Speed Configuration** displays AUTO. For more information, see *speed*.

When you use the **show interfaces extensive** command to view the speed of the interface, the output does not display support for auto-negotiation. However, autonegotiation is supported when the interface speed is configured for 1-Gbps speed.

**NOTE:** You cannot configure the 10-gigabit Ethernet interface, operating with a speed of 1-Gbps, as a member interface of a link aggregation group (LAG).

### ***Port Speed on MIC-MACSEC-20GE***

The MIC-MACSEC-20GE MIC provides 128-bit and 256-bit MACsec encryption on all the twenty 1GE and on the two 10GE ports in the following hardware configuration:

- Installed directly on the MX80 and MX104 routers
- Installed on MPC1, MPC2, MPC3, MPC2E, MPC3E, MPC2E-NG, and MPC3E-NG line cards on the MX240, MX480, and MX960 routers

By default, 128-bit MACsec encryption is supported.

The twenty 1-Gigabit Ethernet SFP ports distribute the ports across PIC0 and PIC1, that are logical PICs on the physical MIC. The two 10-Gigabit Ethernet SFP+ ports are physically located on PIC1. But, the 10-Gigabit interfaces are created by distributing the ports in either of the PICs. For information about the naming conventions for interfaces on MIC-MACSEC-20GE, see *Interface Naming Conventions for MIC-MACSEC-20GE*.

#### **NOTE:**

- When the pic-mode is changed from 1-Gbps to 10-Gbps or vice versa, the Flexible PIC Concentrator (FPC) in MX240, MX480, MX960 routers and the Forwarding Engine Board (FEB) in MX80, MX104 routers undergoes an automatic bounce or reboot.
- When the MIC-MACSEC-20GE is operating in the 10-Gbps mode, all the other 1-Gbps ports are disabled.

### **SEE ALSO**

*Configuring Port Speed on MPC7E (Multi-Rate) to Enable Different Port Speeds*

*Configuring Port Speed on MIC-MRATE to Enable Different Port Speeds*

[Junos Continuity Software User Guide \(Junos OS Release 14.1R4 and Later Releases\)](#)

[Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription](#)

*Port Speed*

## Guidelines for Configuring Port Speed on Routing Devices

### IN THIS SECTION

- [Guidelines for Configuring Port Speed for MIC-MRATE MIC and MPC7E-MRATE MPC | 374](#)
- [Guidelines for Configuring Port Speed for JNP10K-LC2101 | 375](#)
- [Guidelines for Configuring Port Speed for MPC11E | 376](#)

This topic describes the guidelines to consider when configuring rate selectability at the port level or the PIC or MIC level.

### *Guidelines for Configuring Port Speed for MIC-MRATE MIC and MPC7E-MRATE MPC*

This topic describes the guidelines to consider when configuring rate selectability at the port level or the PIC level for MIC-MRATE MIC and MPC7E-MRATE MPC:

- If rate selectability is not configured, all ports of the MIC-MRATE MIC and MPC7E-MRATE MPC operate as four 10-Gigabit Ethernet interfaces by default. Therefore, when booting the MPC:
  - If rate selectability is not configured or if invalid port speeds are configured, an alarm is generated to indicate that the configuration is invalid. All the ports operate as four 10-Gigabit Ethernet interfaces.
  - If valid port speeds are configured, the PIC and MIC operate at the configured speed.
- When you change an existing port speed configuration at the port level, you must reset the MPC7E-MRATE PIC for the configuration to take effect.

Similarly, when you change an existing port speed configuration at the port level for MPC8E or MPC9E, you must reset the MIC for the configuration to take effect. You can use the **request chassis mic mic-slot *mic-slot-number* fpc-slot *fpc-slot-number* (offline | online)** command to reset the MIC and apply your configuration changes.

An alarm is generated indicating the change in port speed configuration.

- When you change an existing port speed configuration with an *invalid* port speed configuration, an alarm is generated indicating that the port speed configuration is invalid. For example, on the MPC7E-MRATE, if you configure the port speed of port 3 as 100 Gbps, it is an invalid configuration. MPC7E -MRATE supports 100 Gbps only on ports 2 and 5. The MPC continues to operate using the existing port speed configuration or the default port speed.
- You cannot configure rate selectability at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations.
- When you configure rate selectability at the port level, only the configured ports are enabled. Other ports are disabled.

### ***Guidelines for Configuring Port Speed for JNP10K-LC2101***

This topic describes the guidelines to consider when configuring rate selectability at the port level or the PIC level for JNP10K-LC2101:

- Each port on the JNP10K-LC2101 MPC supports speeds of 10 Gbps (using breakout cables), 40 Gbps, and 100 Gbps. However, JNP10K-LC2101 MPC does not support bandwidth oversubscription. So, when you configure the ports on all PICs, ensure that the demand on each Packet Forwarding Engine is less than or equal to its forwarding capacity. The default port speed for all PICs is 10G.
- When you change an existing port speed configuration at the port level, you must reset the PIC for the configuration to take effect. When you change an existing port speed configuration at the PIC level, the JNP10K-LC2101 automatically resets the PIC.
- When you change the number of active ports using the **number-of-ports** command, you must reset the PIC for the configuration to take effect. 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.

**NOTE:** You cannot configure the number of active ports at the port level. If you attempt to configure the number of active ports at the port level, an error message is displayed.

- You cannot configure rate selectability at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations.
- When you change an existing port speed configuration with an *invalid* port speed configuration, an alarm is generated indicating that the port speed configuration is invalid. The MPC continues to operate using the existing port speed configuration or the default port speed.

### *Guidelines for Configuring Port Speed for MPC11E*

This topic describes the guidelines to consider when configuring rate selectability at the port level or the PIC level for MPC11E:

- If you do not configure rate selectability at the PIC level using the **pic-mode** option, then the default port speed is 100 Gbps.
- Ports 1 through 4 on each PIC of the MPC11E MPC supports speeds of 100 Gbps.
- On MPC11E, you cannot configure the number of active ports or the number of channelized-interfaces to be created on a port. The **number-of-ports** and **number-of-sub-ports** statements are not supported.
- You cannot configure rate selectability at the PIC level and the port level simultaneously. Error messages are displayed when you attempt to commit the configuration.
- When you configure rate selectability at the port level, only the configured ports are created in that PIC. Other ports are not created. When you change the port configuration at the port level, the interfaces corresponding to the affected port are deleted and then re-created.
- When you change an existing port speed configuration with an *invalid* port speed configuration, an alarm is generated indicating that the port speed configuration is invalid. The MPC continues to operate using the existing port speed configuration or the default port speed.

### SEE ALSO

| *Configuring Port Speed on MPC7E (Multi-Rate) to Enable Different Port Speeds*

### MX10003 MPC Port Speed Overview

#### IN THIS SECTION

- [Invalid Port Configuration | 382](#)

MX10003 MPC supports a Multi-Rate 12xQSFP28 Ethernet Modular Interface Card (MIC) and a fixed-port PIC (6xQSFP). The MX10003 Packet Forwarding Engine has 6x40GE QSFP ports on the fixed-port PIC and 12x100GE QSFP28 ports on the MIC. For more information see [MX10003 MPC \(Multi-Rate\)](#). Rate selectability enables you to configure the port speed either at the port level or at the MIC level. To configure all ports to operate at the same speed, you configure rate selectability at the MIC or

PIC level. For more information see *Configuring Port Speed on MX10003 MPC to Enable Different Port Speeds*. To configure different port speeds for each port, you configure rate selectability at the port level, in which case only the ports that are configured are enabled. For more information see *Configuring Port Speed on MX10003 MPC to Enable Different Port Speeds*.

The ports on the MX10003 MPC are called rate-selectable or multirate ports as they support multiple port speeds. You can choose to configure all supported ports of the fixed port PIC or the MIC to operate at the same speed or configure all the ports at different supported speeds. However, all the PIC or MIC ports do not support all the port speeds. For example, you can choose to configure:

- A port in 4x10GE mode using QSFP-4x10GE optics and 4x10GE breakout cables.
- A port in 40GE mode using QSFP optics.
- A port in 100GE mode using QSFP28 optics.
- A port in 1GE mode (for the ports that is already operating in 10GE mode only) using QSFP-4x10GE optics on fixed PIC and non-MacSEC MIC.

**NOTE:**

- You can use the [port-checker tool](#) to check whether the combination of ports you want to use is valid or not.
- You can use the [Hardware Compatibility Tool](#) to find information about the pluggable transceivers supported on MX10003 router.

The MX10003 MPC supports three Packet Forwarding Engines. The forwarding capacity of each Packet Forwarding Engine is 400Gbps which cannot be oversubscribed.

The MIC supports 12 ports. Each Packet Forwarding Engine is mapped to 4 ports of the MIC. Port 0 through port 3 are mapped to PFE0, port 4 through port 7 are mapped to PFE1, and port 8 through port 11 are mapped to PFE2. The fixed-port PIC supports 6 ports. Each Packet Forwarding Engine is mapped to two ports of the fixed-port PIC. Port 0 and port 1 are mapped to PFE0, port 2 and port 3 are mapped to PFE1, and port 4 and port 5 are mapped to PFE2. You can use the command **show chassis pic fpc-slot slot-number pic-slot slot-number** to display Packet Forwarding Engine mapping information and port speed information.

[Table 43](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 74: Supported Port Speed of MX10003 MPC**

| PIC                    | Port Number | Port Speed Supported                                                                                                                                                                            |
|------------------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (Fixed-port PIC) | 0–5         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br><b>NOTE:</b> You can configure one or all 10G port operating in 4X10-Gigabit Ethernet mode to operate in 1-Gigabit Ethernet mode.               |
| PIC 1 (Multi-Rate MIC) | 0–11        | 100-Gigabit Ethernet<br>40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br><b>NOTE:</b> On non-MACsec MIC, you can configure one or all the 4X10-Gigabit Ethernet port to 1-Gigabit Ethernet mode. |

Starting with Junos OS Release 18.1R1, the non-MACsec MIC on the MX10003 routers support 1-Gigabit Ethernet mode also on 10-Gigabit Ethernet mode ports.

Each of the 100-Gigabit Ethernet or 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 can also use 4x10GE LR breakout optics (QSFP-4X10GE-LR) at the MX10003 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. Refer to [Hardware Compatibility Tool](#) for the list of pluggable transceivers supported on MX10003 router.

On MX10003 routers, when the port operates in 10-Gbps speed, you can change the operating speed to 1Gbps using the configuration **speed 1G** as follows:

```
set interfaces interface-name gigether-options speed 1g
```

Refer *speed (Gigabit Ethernet interface)* for more details.

Once you commit this configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed, but the **show interface** command displays for the field **Physical interface** (that is, the interface name prefix) as **XE /\_/** and the **Speed Configuration** (that is, operating port speed) as **1GE**. On fixed-port PIC and non-MACsec MIC, you can configure one or all 10-Gbps port operating in 4X10-Gbps speed to operate in 1-Gbps speed.

1-Gbps speed is only supported in non-autonegotiation mode.

**NOTE:**

- Any interface operating at 10-Gigabit Ethernet mode can be independently converted to 1-Gigabit Ethernet mode. For example, in multi-rate connections through split cables, when one of the ports operates at 1GE mode, the other three ports can still be configured in 1GE or 10GE mode.
- The MACsec MIC does not support 1-Gbps speed.
- The rate selectability at PIC level and port level does not support 1-Gbps speed. But you can configure the port configured at 10-Gbps speed to operate at 1-Gbps speed using the *speed (Gigabit Ethernet interface)* configuration statement at Gigabit Ethernet interface level.
- The 1-Gbps operation mode is only supported in non-autonegotiation mode.
- ISSU is not supported for the interfaces that are configured with 1-Gigabit Ethernet mode. If ISSU upgrade is carried out in 1-Gigabit Ethernet mode, then the behavior is unexpected and traffic loss can be expected. Refer *request vmhost software in-service-upgrade* for more details.

To view the speed configured for the interface, execute the *show interfaces extensive* command. The **Speed Configuration** output parameter in the command output indicates the current operation speed of the interface. If the interface is configured with 1-Gbps speed, then **Speed Configuration** displays **1G**; if the interface is configured with 10-Gbps speed, **Speed Configuration** displays **AUTO**.

For example:

```
user@host>show interfaces xe-0/1/11:0 extensive
Physical interface: xe-0/1/11:0, Enabled, Physical link is Up
Interface index: 284, SNMP ifIndex: 609, Generation: 383
Link-level type: Ethernet, MTU: 9192, MRU: 9200, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
Loopback: None, Source filtering: Disabled, Flow control: Enabled,
Speed Configuration: 1G
...
```

In this example, the **Speed Configuration** output parameter displays 1G, which means the operation speed of xe-0/1/11:0 interface is 1-Gbps speed.

MX10003 MPC has an aggregate forwarding capacity of 1.2 Tbps and a forwarding capacity of 400 Gbps on each Packet Forwarding Engine. Oversubscription of Packet Forwarding Engine capacity is not supported. The demand on each Packet Forwarding Engine must be less than or equal to its forwarding

capacity. For more information see, [Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription on MX10003 MPC](#). For instance, for MX10003 MPC, the demand on each Packet Forwarding Engine must be less than or equal to 400 Gbps.

For example, on the fixed-port PIC, if you configure the port speed on one ports as 40 Gbps or on two port as 40 Gbps, then you can configure the ports on the MIC in one of the following ways:

- Three 100-Gigabit Ethernet interfaces
- Two 100-Gigabit Ethernet and two 40-Gigabit Ethernet interfaces
- Two 100-Gigabit Ethernet and eight 10-Gigabit Ethernet interfaces

The same rule is applicable to all Packet Forwarding Engines independently.

**NOTE:** Only the Interface that is already operating at 10GE mode can be configured to operate at 1GE mode using `speed (Gigabit Ethernet interface)` configuration statement as follows:  
**set interfaces *interface-name* gigheter-options speed 1g**

[Table 44](#) summarizes the port mode configuration at the Packet Forwarding Engine level.

**Table 75: PFE Based Port Mode Configuration**

| Port Speed configuration on PIC1(Gbps) |       |       |       | Port speed configuration on PIC0(Gbps) |       |
|----------------------------------------|-------|-------|-------|----------------------------------------|-------|
| 100                                    | 100   | 100   | 100   | 0                                      | 0     |
| 100                                    | 100   | 100   | 10/40 | 0                                      | 0     |
| 100                                    | 100   | 100   | 0     | 10/40                                  | 10/40 |
| 100                                    | 100   | 10/40 | 10/40 | 10/40                                  | 10/40 |
| 100                                    | 10/40 | 10/40 | 10/40 | 10/40                                  | 10/40 |
| 10/40                                  | 10/40 | 10/40 | 0     | 10/40                                  | 10/40 |
| 10/40                                  | 10/40 | 10/40 | 10/40 | 0                                      | 0     |

Table 45 summarizes the PIC mode configuration.

**Table 76: PIC Mode Configuration**

| Port Speed configuration on PIC1(Gbps) | Port speed configuration on PIC0(Gbps)   |
|----------------------------------------|------------------------------------------|
| 100                                    | 0<br>Configure the number of ports to 0. |
| 10                                     | 10                                       |
| 40                                     | 40                                       |

Note the following caveats while configuring rate selectability on the MX10003 MPC:

- By default, the MX10003 router comes up with the PIC mode where all the interface operates at the same speed of 10-Gbps. That is, by default, both the PICs (PIC 0 and PIC 1) operate at 10-Gbps speed. To use different port speeds, you must first switch to the port mode and then change the default speed.

To change the default speed, you must select a port and configure a different port speed on it and reset both the PICs for the configuration to take effect. For example, select 40GE or 100GE on PIC 1 and 10GE on PIC 0. For this configuration to take effect, you must reset both PICs.

- Regardless of the line card— MIC (PIC1) or fixed-port PIC (PIC0) installed —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.
- The port speed configuration on the fixed-port PIC and the MIC must be homogenous. However, at port level you can configure port speeds in heterogeneous mode. For more information, see *Configuring Rate Selectability on MX10003 MPC at Port Level*.

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.

- When you configure rate selectability at the port level, only the configured ports are active. Other ports are disabled.

- 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 the ports which 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.
- You cannot configure rate selectability at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations.
- When you change an existing port speed configuration at the port level, you must reset the PIC for the configuration to take effect. When you change an existing port speed configuration at the PIC level, the MPC automatically resets the PIC.

### *Invalid Port Configuration*

You cannot configure the ports which 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.

### SEE ALSO

*Configuring Port Speed on MX10003 MPC to Enable Different Port Speeds  
speed (Gigabit Ethernet interface)*

### MX204 Router Port Speed Overview

#### IN THIS SECTION

- [User-Configurable Port Speed of MX204 Routers | 386](#)
- [Maximum number of 10/40/100GE ports Configurable at PIC and Port Mode | 387](#)
- [Port Configuration - PIC Level | 388](#)

The maximum amount of data that can be transmitted through a port at any given second either by a network device or by a component of the network device (such as a line card) is known as the port speed. Port speed is measured in kilobits per second (Kbps), gigabits per second (Gbps), and terabits per second (Tbps). If a port can be configured to support both single and multiple speeds, the port is known as a rate-selectable port. Because the port is part of a network device (router or switch) or a network component (such as MPC, MIC) the component is known as a rate-selectable component. Rate selectability enables you to configure different port speeds at the port level or at the PIC level.

The MX204 has four rate-selectable ports (referred to as PIC 0 ports) that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet port, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The MX204 also has eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports). On PIC 0 and PIC 1, you can configure the 10-Gigabit Ethernet port(s) to operate in 10-Gigabit Ethernet mode (using *speed (Gigabit Ethernet interface)* command). The four rate-selectable ports supports QSFP28/QSFP+ transceivers, whereas the eight 10-Gigabit Ethernet ports supports SFP+ transceivers. Knowing the exact details of the port speeds for the PICs helps you to choose the speeds to configure on the ports or on the PICs. You can view the port speeds of the PIC by executing *show chassis pic* command. For more information, see [MX204 Router Overview](#) and [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router](#).

#### NOTE:

- By default, the MX204 router comes up with the PIC mode where all the interface operates at the same speed of 10-Gbps. that is, by default, both the PICs (PIC 0 and PIC 1) 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 port mode to PIC mode, the PIC restarts automatically.

To change the default speed, you must select a port and configure a different port speed on it and reset both the PICs for the configuration to take effect. For example, select 40GE or 100GE on PIC 0 and 10GE on PIC 1. For this configuration to take effect, you must reset both PICs.

- Not all port combinations will work. So, it is recommended to use the [port-checker tool](#) to check whether the combination of ports you want to use is valid or not.
- You can use the [Hardware Compatibility Tool](#) to find information about the pluggable transceivers supported on MX204 router.

The MX204 router supports two types of rate selectability configuration options:

- **PIC Level Configuration:** To configure all ports to operate at the same speed, you configure rate selectability at the PIC level.

- **Port Level Configuration:** To configure different port speeds for each port, you configure rate selectability at the port level, in which case only the ports that are configured are enabled.

To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the *pic-mode* statement and specify the port speed. To configure different port speeds for each port, configure rate selectability at the port level, in which case only the ports that are configured are enabled. To configure rate selectability at the port level, use the *speed* statement to specify the speed of individual ports.

The examples below show the sample CLI command output of the port speed capability of the 4-port PIC 0 with QSFP+ transceivers and the 8-port PIC 1 with SFP+ transceivers on the MX204 router.

```

user@host> show chassis pic fpc-slot 0 pic-slot 0
...
Port Speed Information:

 Port Capable Port Speeds
 0 4x10GE, 40GE, 100GE
 1 4x10GE, 40GE, 100GE
 2 4x10GE, 40GE, 100GE
 3 4x10GE, 40GE, 100GE
...
user@host> show chassis pic fpc-slot 0 pic-slot 1
...
Port Speed Information:

 Port Capable Port Speeds
 0 10GE
 1 10GE
 2 10GE
 3 10GE
 4 10GE
 5 10GE
 6 10GE
 7 10GE
...

```

[Table 46](#) summarizes the rate selectability of the MX204 routers.

Table 77: Rate Selectability of MX204 Routers

| PIC   | Port Number | Port Speed Supported                                                                                                                                                                                                                                 |
|-------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 | 0-3         | 100-Gigabit Ethernet<br>40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br><b>NOTE:</b> <ul style="list-style-type: none"> <li>• Default port speed is 4x10 Gigabit Ethernet.</li> <li>• Supports 1-Gbps speed on 10 Gigabit Ethernet ports.</li> </ul> |
| PIC 1 | 0-7         | 10 Gigabit Ethernet<br><b>NOTE:</b> Supports 1-Gbps speed on 10 Gigabit Ethernet ports.                                                                                                                                                              |

Starting with Junos OS Release 18.1R1, the 10-Gbps port can operate in 1-Gbps mode also.

Each of the four 100-Gigabit Ethernet or 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 can also use 4x10GE LR breakout optics (QSFPP-4X10GE-LR) at the MX204 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. Refer to [Hardware Compatibility Tool](#) for the list of pluggable transceivers supported on MX204 router.

MX204 router also support 1-Gigabit Ethernet port on the fixed 10-Gigabit Ethernet SFPP ports with 1GE SFPs in it.

On MX204 routers, when the port is operating in 10-Gbps speed, you can change the operating speed to 1Gbps using the configuration statement **Speed 1G** as follows:

```
set interfaces interface-name gigether-options speed 1g
```

Refer *speed (Gigabit Ethernet interface)* for more details.

Once you commit this configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed, but the **show interface** command displays for the field **Physical interface** (that is, the interface name prefix) as **XE /\_/** and the **Speed Configuration** (that is, operating port speed) as **1GE**.

On MRATE PIC, each channel per port can be configured individually as 1-Gigabit Ethernet port.

**NOTE:**

- The interface name prefix must be xe.
- The rate selectability at PIC level and port level does not support 1-Gbps speed. But you can configure the port configured at 10-Gbps speed to operate at 1-Gbps speed using the *speed (Gigabit Ethernet interface)* configuration statement at Gigabit Ethernet interface level.
- The 1-Gbps SFP port supports auto-negotiation. You can configure auto-negotiation by using the command **set interfaces *interface-name* gigheter-options auto-negotiation**. For more information, see *auto-negotiation*.

To view the speed configured for the interface, execute the *show interfaces extensive* command. The **Speed Configuration** output parameter in the command output indicates the current operation speed of the interface. If the interface is configured with 1-Gbps speed, then **Speed Configuration** displays 1G; if the interface is configured with 10-Gbps speed, Speed Configuration displays AUTO.

For example:

```
user@host>show interfaces xe-0/1/11:0 extensive
Physical interface: xe-0/1/11:0, Enabled, Physical link is Up
Interface index: 284, SNMP ifIndex: 609, Generation: 383
Link-level type: Ethernet, MTU: 9192, MRU: 9200, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
Loopback: None, Source filtering: Disabled, Flow control: Enabled,
Speed Configuration: 1G
...
```

In this example, the **Speed Configuration** output parameter displays 1G, which means the operation speed of xe-0/1/11:0 interface is 1-Gbps speed.

### ***User-Configurable Port Speed of MX204 Routers***

You can also configure rate selectability on MX204 routers.

[Table 47](#) summarizes the user-configurable rate selectability of MX204 routers.

**Table 78: Configurable Rate Selectability of MX204 Router**

| Port Speed Configuration on PIC 0 (Gbps) | Port Speed Configuration on PIC 1 (Gbps)        |
|------------------------------------------|-------------------------------------------------|
| 100                                      | 0<br>Configure the number of active ports to 0. |
| 10                                       | 10                                              |
| 40                                       | 0<br>Configure the number of active ports to 0. |

Only the Interface that is already operating at 10GE mode can be configured to operate at 1GE mode using *speed (Gigabit Ethernet interface)* configuration statement as follows:

**set interfaces *interface-name* gigether-options speed *1g***

**NOTE:** 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.

#### ***Maximum number of 10/40/100GE ports Configurable at PIC and Port Mode***

Following table summarizes the maximum number of 10/40/100 Gigabit Ethernet ports per PIC configurable at PIC and port levels:

**Table 79: Maximum number of 10/40/100 Gigabit Ethernet ports Configurable at PIC and Port Level**

| Maximum Ports               | Maximum Ports configurable at PIC Mode (on both PIC0 and PIC1) | Maximum Ports Configurable at Port Mode (on both PIC0 and PIC1) |
|-----------------------------|----------------------------------------------------------------|-----------------------------------------------------------------|
| 10/1 Gigabit Ethernet Ports | 24<br>Which means 16 ports from PIC 0 and 8 Ports from PIC 1.  | 20<br>Which means 12 ports from PIC 0 and 8 Ports from PIC 1.   |

**Table 79: Maximum number of 10/40/100 Gigabit Ethernet ports Configurable at PIC and Port Level (Continued)**

| Maximum Ports              | Maximum Ports configurable at PIC Mode (on both PIC0 and PIC1)     | Maximum Ports Configurable at Port Mode (on both PIC0 and PIC1) |
|----------------------------|--------------------------------------------------------------------|-----------------------------------------------------------------|
| 40 Gigabit Ethernet Ports  | 4<br>Only 4 ports from PIC 0 as PIC 1 supports only 10 Gbps Speed. | 4                                                               |
| 100 Gigabit Ethernet Ports | 4<br>Only 4 ports from PIC 0 as PIC 1 supports only 10 Gbps Speed. | 4                                                               |

#### **Port Configuration - PIC Level**

On PIC 0, if each of the four ports is configured to operate at 100-Gbps speed, then you must configure all the 8 ports at PIC 1 to 0 (using *number-of-ports* statement). On PIC 0, if ports 0, 1, and 2 are set to 100-Gbps, and port 3 is set to 10-Gbps or 40-Gbps, then you should configure all the 8 ports at PIC 1 to 0 (using **number-of-port** statement), and so on as listed in the below table.

The following table only lists few valid combination of port speeds on PIC 0 and PIC1 of MX204 router. You are not limited to configure only the below mentioned example port configurations. For more valid port configuration values, refer [port-checker tool](#).

**Table 80: Port Configuration at PIC Level in MX204 Routers**

| Port Mode |     |     |     |                                                 |
|-----------|-----|-----|-----|-------------------------------------------------|
| PIC 0     |     |     |     | PIC 1                                           |
| 100       | 100 | 100 | 100 | 0<br>Configure the number of active ports to 0. |

Table 80: Port Configuration at PIC Level in MX204 Routers (Continued)

| Port Mode |       |       |       |                                                 |    |    |    |    |    |    |    |    |
|-----------|-------|-------|-------|-------------------------------------------------|----|----|----|----|----|----|----|----|
| PIC 0     |       |       |       | PIC 1                                           |    |    |    |    |    |    |    |    |
| 100       | 100   | 100   | 10/40 | 0<br>Configure the number of active ports to 0. |    |    |    |    |    |    |    |    |
| 100       | 100   | 10/40 | 10/40 | 10                                              | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 10 |
| 100       | 10/40 | 10/40 | 10/40 | 10                                              | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 10 |
| 10/40     | 10/40 | 10/40 | 10/40 | 0<br>Configure the number of ports to 0.        |    |    |    |    |    |    |    |    |

| PIC Mode                                |     |     |     |                                                                |  |  |  |  |  |  |  |  |
|-----------------------------------------|-----|-----|-----|----------------------------------------------------------------|--|--|--|--|--|--|--|--|
| PIC 0 (with four rate-selectable ports) |     |     |     | PIC 1 (with eight 10-Gigabit Ethernet ports)                   |  |  |  |  |  |  |  |  |
| 100                                     | 100 | 100 | 100 | 0<br>Configure the number of active ports to 0.                |  |  |  |  |  |  |  |  |
| 40                                      | 40  | 40  | 40  | 0<br>Configure the number of active ports to 0.                |  |  |  |  |  |  |  |  |
| 10                                      | 10  | 10  | 10  | 10<br>Configure all the eight 10-Gigabit Ethernet ports to 10. |  |  |  |  |  |  |  |  |

**NOTE:** Only the Interface that is already operating at 10GE mode can be configured to operate at 1GE mode using *speed (Gigabit Ethernet interface)* configuration statement as follows:

**set interfaces *interface-name* gigether-options speed *1g***

## SEE ALSO

*Configuring Port Speed on MX204 to Enable Different Port Speeds*

*speed*

*speed (Gigabit Ethernet interface)*

*show chassis pic*

*number-of-ports*

*pic-mode*

## PTX10003 Router Port Speed Overview

PTX10003 Packet Transport Routers feature flexible interface configuration options with universal multi-rate double-density Quad Small Form-factor Pluggable (QSFP-DD) optics. The PTX10003-80C port panel has 40 physical ports and the PTX10003-160C port panel has 80 physical ports. The physical ports are in groups of five QSFP-DD ports. You can configure different data rates for each port group as long as the specified guidelines are met. Any port can be used as a 100-Gigabit Ethernet interface, 40-Gigabit Ethernet interface, 25-Gigabit Ethernet interface, or 10-Gigabit Ethernet interface. You choose the speed by plugging in the appropriate transceiver.

**NOTE:** The center port in each port group (port 2 and port 7) cannot support 1x200 Gbps. To configure a 200 Gbps data rate 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.

You can channelize the Gigabit Ethernet interfaces on PTX10003 routers to create multiple independent Gigabit Ethernet interfaces and then use breakout cables to connect the channelized ports to other servers, storage devices, and routers. Here's the allowable channelization configurations for the optical transceivers supported by the PTX10003:

**Table 81: Channelization Configuration on PTX10003**

| QSFP Transceiver | Native Port Speeds       | Channelization Options |
|------------------|--------------------------|------------------------|
| QSFP56-DD        | 1x400G Gbps              | 4x100G Gbps            |
| QSFP28-DD        | 1x200 Gbps<br>2x100 Gbps | 8x25 Gbps              |
| QSFP28           | 1x100 Gbps               | 4x25 Gbps              |
| QSFP+            | 1x40 Gbps<br>4x10 Gbps   | 4x10 Gbps              |

**NOTE:** For more details about supported transceivers and cable specifications, see the [PTX10003 Packet Transport Router Hardware Guide](#).

By default, all PTX10003 QSFP-DD interfaces are configured for a data rate of 2x100 Gbps. The interface names appear in the `et-fpc/pic/port:channel` format.

The port speed can be configured at the PIC-level by using the `set chassis fpc slot-number pic slot-number pic-mode pic-mode` command. The `pic-mode` statement can take values 10G, 40G, or 100G to operate all ports in 4x10G, 1x40G, or 1x100G.

To configure the port speed or channelize a port:

1. Issue the following command to set the port speed: `set chassis fpc slot-number pic pic-number port port-number number-of-subports [1 / 2 / 4 / 8] speed [10G / 40G / 100G / 200G / 400G]`

For example, to configure the second port in the first port group as a 1x40 Gbps interface, issue the `set chassis fpc 0 pic 0 port 1 number-of-subports 1 speed 40g` command.

2. Type the `commit` command.

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

After you commit this configuration, the second port in PIC 0 will operate at 1x40 Gbps.

**NOTE:** When a 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 there isn't a port speed configured on a valid optical port, the PTX10003 uses a default port speed of 2x100 Gbps. Also, if **number-of-subports** is not configured, a 1x *40G / 100G / 200G / 400G* data rate is assumed. A 1x10G sub-port is not supported.

When a port is channelized, the interface name has a colon followed by the port channel to signify the four separate channels. For example, on a PTX10003 with port 2 on PIC 1 configured as four 25-Gigabit Ethernet ports, the interface names are et-0/1/2:0, et-0/1/2:1, et-0/1/2:2, and et-0/1/2:3.

There is no commit check when you channelize a port or configure the speed of the port.

## SEE ALSO

[number-of-sub-ports | 1042](#)

[speed \(Ethernet\) | 1137](#)

*fpc*

*pic*

## PTX10K-LC1201 Port Speed Overview

The PTX10K-LC1201 line card is a fixed-configuration, rate-selectable line card with 36 built-in ports. The ports on the PTX10K-LC1201 are called rate-selectable or multi-rate ports as they support multiple port speeds. Rate selectability enables you to configure the port speed either at the port level or at the PIC level. To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the port level, use the **pic-mode** statement and specify the port speed. You can choose to configure all ports to operate at the same speed or configure all the ports to operate at different supported speeds. The default port speed is 400Gbps for all ports. Each PTX10K-LC1201 provides a maximum bandwidth of 14.4Tbps.

On the PTX10K-LC1201, you can choose to configure all 36 ports with the following port speeds:

- 4x10 Gbps, 4x25 Gbps, and 2x50 Gbps
- 40 Gbps, 100 Gbps, 200 Gbps, and 400 Gbps

**NOTE:** When you change the speed at the PIC level, the existing interfaces are deleted and new interfaces are created based on the new configuration. When you change the speed of a particular port explicitly by using the **speed** statement, only that port is affected. All other ports in that PIC remain unaffected.

For information about how to configure rate selectability, see *Configuring Port Speed on PTX10K-1201 line card to Enable Different Port Speeds*.

Table 51 summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 82: Port Speed for the PTX10K-LC1201**

| PIC   | Port Number | Port Speed Supported                                                                                                                                                                                     | Optics Supported                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 | 0-35        | 4x10-Gigabit Ethernet<br>1x40-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>2x50-Gigabit Ethernet<br>1x100-Gigabit Ethernet<br>2x100-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>1x400-Gigabit Ethernet | <ul style="list-style-type: none"> <li>• 1x40GE and 4x10GE support using QSFP+</li> <li>• 4x25GE support using QSFP28 25G optics (using breakout cables).</li> <li>• 2x50GE support using QSFP28 50G optics (using breakout cables).</li> <li>• 1x100GE support using QSFP28 100G optics.</li> <li>• 2x100GE support using QSFP28 DD 200G</li> <li>• 4x100GE and 1x400GE support using QSFP56 DD 400G</li> </ul> <p><b>NOTE:</b> By default, all the active ports operate in 400-Gigabit Ethernet mode.</p> |

## SEE ALSO

[PTX10K-LC1201-36CD for PTX10008 Routers](#)

## RELATED DOCUMENTATION

[Preventing Oversubscription Using Active Physical Ports](#)

*Configuring Port Speed on Routing Devices and Line Cards*

## Interface Naming Conventions

### IN THIS SECTION

- [Interface Naming Conventions for MPC7E-MRATE | 394](#)
- [Interface Naming Conventions for MIC-MRATE | 396](#)
- [Interface Naming Conventions for MX10003 MPC | 400](#)
- [Interface Naming Conventions for PTX10K-LC2101 | 403](#)
- [Interface Naming Conventions for MIC-MACSEC-20GE | 405](#)
- [Interface Naming Conventions for PTX10K-LC1201 | 407](#)
- [Interface Mapping and Modulation Format for ACX5448-D | 416](#)

The interface name uniquely identifies an individual network connector in the system. Use the interface name when you configure the interface. Every device follows its own naming convention. Use this topic to understand more about the interface naming conventions for the line cards and devices.

### Interface Naming Conventions for MPC7E-MRATE

MPC7E (MPC7E-MRATE) is a fixed-configuration MPC and contains two built-in PICs, PIC 0 and PIC 1. Each of the six ports of PIC 0 and PIC 1 support multiple port speeds of 100 Gbps, 40 Gbps, and 10 Gbps and can be configured as 10-Gigabit Ethernet and 40-Gigabit Ethernet interfaces. However, you can configure only ports 2 and 5 on both the PICs as 100-Gigabit Ethernet interfaces.

MPC7E-MRATE has an aggregate forwarding capacity of 480 Gbps and a forwarding capacity of 240 Gbps on each Packet Forwarding Engine. Oversubscription of Packet Forwarding Engine capacity is not supported. The demand on each Packet Forwarding Engine must be less than or equal to its forwarding capacity. For instance, for MPC7E-MRATE, the demand on each Packet Forwarding Engine must be less than or equal to 240 Gbps.

The 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces configured on the MPC7E-MRATE MPC follow the naming convention **et-*fpc-slot/pic-slot/port-number***. The 10-Gigabit Ethernet interfaces configured on the MPC7E-MRATE MPC follow the naming convention **xe-*fpc-slot/pic-slot/port-number.[logical-port-number]***.

For example, **et-0/0/2** indicates either a 40-Gigabit Ethernet or a 100-Gigabit Ethernet interface configured on port 2 of PIC 0 of the MPC7E-MRATE MPC that is installed in the MPC slot 0. **xe-0/0/1:3** indicates a 10-Gigabit Ethernet interface configured on logical port 3 of physical port 1 of the MPC7E-MRATE MPC that is installed in the MPC slot 0.

[Table 52](#) lists the naming conventions for interfaces on MPC7E-MRATE for MX240, MX480, MX960, MX2010, and MX2020 routers.

**Table 83: Interface Naming Convention for MPC7E-MRATE**

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 0                        | <b>xe-x/0/0:[0-3]</b>         | <b>et-x/0/0</b>               | -                              |
|                          | <b>xe-x/0/1:[0-3]</b>         | <b>et-x/0/1</b>               | -                              |
|                          | <b>xe-x/0/2:[0-3]</b>         | <b>et-x/0/2</b>               | <b>et-x/0/2</b>                |
|                          | <b>xe-x/0/3:[0-3]</b>         | <b>et-x/0/3</b>               | -                              |
|                          | <b>xe-x/0/4:[0-3]</b>         | <b>et-x/0/4</b>               | -                              |
|                          | <b>xe-x/0/5:[0-3]</b>         | <b>et-x/0/5</b>               | <b>et-x/0/5</b>                |
| 1                        | <b>xe-x/1/0:[0-3]</b>         | <b>et-x/1/0</b>               | -                              |
|                          | <b>xe-x/1/1:[0-3]</b>         | <b>et-x/1/1</b>               | -                              |
|                          | <b>xe-x/1/2:[0-3]</b>         | <b>et-x/1/2</b>               | <b>et-x/1/2</b>                |
|                          | <b>xe-x/1/3:[0-3]</b>         | <b>et-x/1/3</b>               | -                              |

**Table 83: Interface Naming Convention for MPC7E-MRATE (Continued)**

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
|                          | <i>xe-x/1/4:[0-3]</i>         | <i>et-x/1/4</i>               | -                              |
|                          | <i>xe-x/1/5:[0-3]</i>         | <i>et-x/1/5</i>               | <i>et-x/1/5</i>                |

**SEE ALSO**

*Configuring Port Speed on MPC7E (Multi-Rate) to Enable Different Port Speeds*

**Interface Naming Conventions for MIC-MRATE**

MIC-MRATE consists of twelve ports that support multiple port speeds of 100 Gbps, 40 Gbps, and 10 Gbps. MIC-MRATE is supported on MPC8E (MX2K-MPC8E) and MPC9E (MX2K-MPC9E) on MX2000 line of routers.

Starting with Junos OS Release 17.3R1, MIC-MRATE is supported on MX10003 MPC on MX10003 routers.

**NOTE:** By default, the MIC-MRATE ports are configured as 10-Gigabit Ethernet ports.

MPC8E has a forwarding capacity of 240 Gbps for each Packet Forwarding Engine. In Junos OS Release 16.1R1 and later, you can upgrade MPC8E to provide an increased bandwidth of 1600 Gbps (1.6 Tbps), by using an add-on license. After you configure the `bandwidth 1.6T` statement, MPC8E provides an increased bandwidth of 1.6 Tbps. The forwarding capacity is increased to 400 Gbps for each Packet Forwarding Engine.

MPC9E has a forwarding capacity of 400 Gbps for each Packet Forwarding Engine. Packet Forwarding Engine oversubscription is not supported. So, demand on each Packet Forwarding Engine should be less than or equal to its forwarding capacity. For MPC8E, demand on each Packet Forwarding Engine should be less than or equal to 240 Gbps and for MPC9E, demand on each Packet Forwarding Engine should be less than or equal to 400 Gbps.

**NOTE:** On MPC8E with MIC-MRATE, you can configure four ports as 100-Gigabit Ethernet interfaces. On MPC9E with MIC-MRATE and on MPC8E configured to operate at 1.6 Tbps by using an add-on license, you can configure eight ports as 100-Gigabit Ethernet interfaces.

The 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces configured on the MIC-MRATE MIC follow the naming convention **et-*fpc-slot/pic-slot/port-number***. The 10-Gigabit Ethernet interfaces configured on the MIC-MRATE MIC follow the naming convention **xe-*fpc-slot/pic-slot/port-number***.**[*logical-port-number*]**.

For example, **xe-0/0/1:3** indicates a 10-Gigabit Ethernet interface configured on logical port 3 of physical port 1 of the MIC-MRATE MIC that is installed in the MPC slot 0. The interface name **et-0/0/2** indicates either a 40-Gigabit Ethernet interface or a 100-Gigabit Ethernet interface configured on port 2 of MIC-MRATE MIC that is installed in the MPC slot 0.

[Table 53](#) lists the naming conventions used for interfaces on MIC-MRATE when installed on slot 0 of MPC8E and MPC9E. [Table 54](#) lists the naming conventions used for interfaces on MIC-MRATE when installed on slot 1 of MPC8E and MPC9E. MPC8E and MPC9E support two MIC-MRATE MICs each.

**NOTE:** The *x* in **et-*x*/0/0** and **xe-*x*/0/0:[0-3]** refers to the MPC slot number.

**Table 84: Interface Naming Convention for MIC-MRATE Installed on Slot 0 of MPC8E and MPC9E**

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 0                        | <b>xe-<i>x</i>/0/0:[0-3]</b>  | <b>et-<i>x</i>/0/0</b>        | <b>et-<i>x</i>/0/0</b>         |
|                          | <b>xe-<i>x</i>/0/1:[0-3]</b>  | <b>et-<i>x</i>/0/1</b>        | <b>et-<i>x</i>/0/1</b>         |
|                          | <b>xe-<i>x</i>/0/2:[0-3]</b>  | <b>et-<i>x</i>/0/2</b>        | <b>et-<i>x</i>/0/2</b>         |
|                          | <b>xe-<i>x</i>/0/3:[0-3]</b>  | <b>et-<i>x</i>/0/3</b>        | <b>et-<i>x</i>/0/3</b>         |
|                          | <b>xe-<i>x</i>/0/4:[0-3]</b>  | <b>et-<i>x</i>/0/4</b>        | -                              |

**Table 84: Interface Naming Convention for MIC-MRATE Installed on Slot 0 of MPC8E and MPC9E**  
(Continued)

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
|                          | <i>xe-x/0/5:[0-3]</i>         | <b>et-x/0/5</b>               | -                              |
| 1                        | <i>xe-x/0/6:[0-3]</i>         | <b>et-x/0/6</b>               | <b>et-x/0/6</b>                |
|                          | <i>xe-x/0/7:[0-3]</i>         | <b>et-x/0/7</b>               | <b>et-x/0/7</b>                |
|                          | <i>xe-x/0/8:[0-3]</i>         | <b>et-x/0/8</b>               | <b>et-x/0/8</b>                |
|                          | <i>xe-x/0/9:[0-3]</i>         | <b>et-x/0/9</b>               | <b>et-x/0/9</b>                |
|                          | <i>xe-x/0/10:[0-3]</i>        | <b>et-x/0/10</b>              | -                              |
|                          | <i>xe-x/0/11:[0-3]</i>        | <b>et-x/0/11</b>              | -                              |

**Table 85: Interface Naming Convention for MIC-MRATE Installed on Slot 1 of MPC8E and MPC9E**

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 2                        | <i>xe-x/1/0:[0-3]</i>         | <b>et-x/1/0</b>               | <b>et-x/1/0</b>                |
|                          | <i>xe-x/1/1:[0-3]</i>         | <b>et-x/1/1</b>               | <b>et-x/1/1</b>                |
|                          | <i>xe-x/1/2:[0-3]</i>         | <b>et-x/1/2</b>               | <b>et-x/1/2</b>                |
|                          | <i>xe-x/1/3:[0-3]</i>         | <b>et-x/1/3</b>               | <b>et-x/1/3</b>                |
|                          | <i>xe-x/1/4:[0-3]</i>         | <b>et-x/1/4</b>               | -                              |

**Table 85: Interface Naming Convention for MIC-MRATE Installed on Slot 1 of MPC8E and MPC9E**  
(Continued)

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
|                          | <i>xe-x/1/5:[0-3]</i>         | <b>et-x/1/5</b>               | -                              |
| 3                        | <i>xe-x/1/6:[0-3]</i>         | <b>et-x/1/6</b>               | <b>et-x/1/6</b>                |
|                          | <i>xe-x/1/7:[0-3]</i>         | <b>et-x/1/7</b>               | <b>et-x/1/7</b>                |
|                          | <i>xe-x/1/8:[0-3]</i>         | <b>et-x/1/8</b>               | <b>et-x/1/8</b>                |
|                          | <i>xe-x/1/9:[0-3]</i>         | <b>et-x/1/9</b>               | <b>et-x/1/9</b>                |
|                          | <i>xe-x/1/10:[0-3]</i>        | <b>et-x/1/10</b>              | -                              |
|                          | <i>xe-x/1/11:[0-3]</i>        | <b>et-x/1/11</b>              | -                              |

Table 55 lists the naming conventions used for interfaces on MIC-MRATE when installed on slot 1 of MX10003 MPC.

**Table 86: Interface Naming Convention for MIC-MRATE Installed on Slot 1 of MX10003MPC**

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 0                        | <i>xe-x/1/0:[0-3]</i>         | <b>et-x/1/0</b>               | <b>et-x/1/0</b>                |
|                          | <i>xe-x/1/1:[0-3]</i>         | <b>et-x/1/1</b>               | <b>et-x/1/1</b>                |
|                          | <i>xe-x/1/2:[0-3]</i>         | <b>et-x/1/2</b>               | <b>et-x/1/2</b>                |
|                          | <i>xe-x/1/3:[0-3]</i>         | <b>et-x/1/3</b>               | <b>et-x/1/3</b>                |

**Table 86: Interface Naming Convention for MIC-MRATE Installed on Slot 1 of MX10003MPC**  
*(Continued)*

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 1                        | <i>xe-x/1/4:[0-3]</i>         | <b>et-x/1/4</b>               | -                              |
|                          | <i>xe-x/1/5:[0-3]</i>         | <b>et-x/1/5</b>               | -                              |
|                          | <i>xe-x/1/6:[0-3]</i>         | <b>et-x/1/6</b>               | <b>et-x/1/6</b>                |
|                          | <i>xe-x/1/7:[0-3]</i>         | <b>et-x/1/7</b>               | <b>et-x/1/7</b>                |
| 2                        | <i>xe-x/1/8:[0-3]</i>         | <b>et-x/1/8</b>               | <b>et-x/1/8</b>                |
|                          | <i>xe-x/1/9:[0-3]</i>         | <b>et-x/1/9</b>               | <b>et-x/1/9</b>                |
|                          | <i>xe-x/1/10:[0-3]</i>        | <b>et-x/1/10</b>              | -                              |
|                          | <i>xe-x/1/11:[0-3]</i>        | <b>et-x/1/11</b>              | -                              |

## SEE ALSO

*Configuring Port Speed on MIC-MRATE to Enable Different Port Speeds*

## Interface Naming Conventions for MX10003 MPC

The MX10003 MPC supports a Multi-Rate 12xQSFP28 Ethernet MIC (model numbers: JNP-MIC1 and JNP-MIC1-MACSEC) and the fixed-port PIC (6xQSFP).

Each of the 6 ports of the PIC supports 10-Gigabit Ethernet and 40-Gigabit Ethernet interfaces. Each of the 12 ports of the modular MIC supports 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet interfaces. All the ports of the modular MIC can be configured as 100-Gigabit Ethernet interfaces.

The 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces configured on the MX10003 MPC follow the naming convention **et-*fpc-slot/pic-slot/port-number***. The 10-Gigabit Ethernet interfaces follow the naming convention **xe-*fpc-slot/pic-slot/port-number.[logical-port-number]***.

For example, **xe-1/1/1:3** indicates a 10-Gigabit Ethernet interface configured on logical port 3 of physical port 1 of the modular MIC that is installed in the MPC slot 1. The interface name **et-1/1/2** indicates either a 40-Gigabit Ethernet interface or a 100-Gigabit Ethernet interface configured on port 2 of modular MIC that is installed in the MPC slot 1.

**NOTE:** The *x* in **et-*x/0/0*** and **xe-*x/0/0:[0-3]*** refers to the MPC slot number.

[Table 56](#) lists the naming conventions used for interfaces on the fixed-port PIC when installed in slot 0 of the MX10003 MPC. [Table 57](#) lists the naming conventions used for interfaces on the modular MIC when installed in slot 1 of the MPC.

**Table 87: 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                        | <b>xe-<i>x/0/0:[0-3]</i></b>  | <b>et-<i>x/0/0</i></b>        | -                              |
|                          | <b>xe-<i>x/0/1:[0-3]</i></b>  | <b>et-<i>x/0/1</i></b>        | -                              |
| 1                        | <b>xe-<i>x/0/2:[0-3]</i></b>  | <b>et-<i>x/0/2</i></b>        | -                              |
|                          | <b>xe-<i>x/0/3:[0-3]</i></b>  | <b>et-<i>x/0/3</i></b>        | -                              |
| 2                        | <b>xe-<i>x/0/4:[0-3]</i></b>  | <b>et-<i>x/0/4</i></b>        | -                              |
|                          | <b>xe-<i>x/0/5:[0-3]</i></b>  | <b>et-<i>x/0/5</i></b>        | -                              |

Table 88: 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                        | <i>xe-x/1/0:[0-3]</i>         | <b>et-x/1/0</b>               | <b>et-x/1/0</b>                |
|                          | <i>xe-x/1/1:[0-3]</i>         | <b>et-x/1/1</b>               | <b>et-x/1/1</b>                |
|                          | <i>xe-x/1/2:[0-3]</i>         | <b>et-x/1/2</b>               | <b>et-x/1/2</b>                |
|                          | <i>xe-x/1/3:[0-3]</i>         | <b>et-x/1/3</b>               | <b>et-x/1/3</b>                |
| 1                        | <i>xe-x/1/4:[0-3]</i>         | <b>et-x/1/4</b>               | <b>et-x/1/4</b>                |
|                          | <i>xe-x/1/5:[0-3]</i>         | <b>et-x/1/5</b>               | <b>et-x/1/5</b>                |
|                          | <i>xe-x/1/6:[0-3]</i>         | <b>et-x/1/6</b>               | <b>et-x/1/6</b>                |
|                          | <i>xe-x/1/7:[0-3]</i>         | <b>et-x/1/7</b>               | <b>et-x/1/7</b>                |
| 2                        | <i>xe-x/1/8:[0-3]</i>         | <b>et-x/1/8</b>               | <b>et-x/1/8</b>                |
|                          | <i>xe-x/1/9:[0-3]</i>         | <b>et-x/1/9</b>               | <b>et-x/1/9</b>                |
|                          | <i>xe-x/1/10:[0-3]</i>        | <b>et-x/1/10</b>              | <b>et-x/1/10</b>               |
|                          | <i>xe-x/1/11:[0-3]</i>        | <b>et-x/1/11</b>              | <b>et-x/1/11</b>               |

## SEE ALSO

| [MX10003 MPC \(Multi-Rate\)](#)

## Interface Naming Conventions for PTX10K-LC2101

PTX10K-LC2101 is a fixed-configuration MPC and contains six built-in PICs, PIC 0 to PIC 5. Each PIC supports 4 ports. All ports support multiple port speeds of 100 Gbps, 40 Gbps, and 10 Gbps and can be configured as 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet interfaces.

PTX10K-LC2101 has a forwarding capacity of 240 Gbps for each Packet Forwarding Engine. PTX10K-LC2101 has six Packet Forwarding Engines. In Junos OS Release 18.2R1 and later, you can upgrade PTX10K-LC2101 to provide an increased bandwidth of 2400 Gbps (2.4Tbps), by using an add-on license. After you configure the `bandwidth 2.4T` statement, PTX10K-LC2101 provides an increased bandwidth of 2.4 Tbps. The forwarding capacity is increased to 400 Gbps for each Packet Forwarding Engine. Packet Forwarding Engine oversubscription is not supported. So, demand on each Packet Forwarding Engine should be less than or equal to its forwarding capacity.

The 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces configured on the PTX10K-LC2101 MPC follow the naming convention `et-fpc-slot/pic-slot/port-number`. The 10-Gigabit Ethernet interfaces configured on the PTX10K-LC2101 MPC follow the naming convention `xe-fpc-slot/pic-slot/port-number.[logical-port-number]`.

For example, `xe-0/0/1:3` indicates a 10-Gigabit Ethernet interface configured on logical port 3 of physical port 1 of the PTX10K-LC2101 MPC that is installed in the MPC slot 0. The interface name `et-0/0/2` indicates either a 40-Gigabit Ethernet interface or a 100-Gigabit Ethernet interface configured on port 2 of the PTX10K-LC2101 MPC that is installed in the MPC slot 0.

**NOTE:** Each Packet Forwarding Engine maps to a single built-in PIC on the PTX10K-LC2101.

Table 58 lists the naming conventions used for interfaces on PTX10K-LC2101 for MX10008 routers. MX10008 routers support 8 PTX10K-LC2101 MPCs.

**NOTE:** The *x* in `et-x/0/0` and `xe-x/0/0:[0-3]` refers to the MPC slot number.

**Table 89: Interface Naming Convention for PTX10K-LC2101 MPC**

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 0                        | <code>xe-x/0/0:[0-3]</code>   | <code>et-x/0/0</code>         | <code>et-x/0/0</code>          |

Table 89: Interface Naming Convention for PTX10K-LC2101 MPC (Continued)

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
|                          | <i>xe-x/0/1:[0-3]</i>         | <i>et-x/0/1</i>               | <i>et-x/0/1</i>                |
|                          | <i>xe-x/0/2:[0-3]</i>         | <i>et-x/0/2</i>               | <i>et-x/0/2</i>                |
|                          | <i>xe-x/0/3:[0-3]</i>         | <i>et-x/0/3</i>               | <i>et-x/0/3</i>                |
| 1                        | <i>xe-x/1/0:[0-3]</i>         | <i>et-x/1/0</i>               | <i>et-x/1/0</i>                |
|                          | <i>xe-x/1/1:[0-3]</i>         | <i>et-x/1/1</i>               | <i>et-x/1/1</i>                |
|                          | <i>xe-x/1/2:[0-3]</i>         | <i>et-x/1/2</i>               | <i>et-x/1/2</i>                |
|                          | <i>xe-x/1/3:[0-3]</i>         | <i>et-x/1/3</i>               | <i>et-x/1/3</i>                |
| 2                        | <i>xe-x/2/0:[0-3]</i>         | <i>et-x/2/0</i>               | <i>et-x/2/0</i>                |
|                          | <i>xe-x/2/1:[0-3]</i>         | <i>et-x/2/1</i>               | <i>et-x/2/1</i>                |
|                          | <i>xe-x/2/2:[0-3]</i>         | <i>et-x/2/2</i>               | <i>et-x/2/2</i>                |
|                          | <i>xe-x/2/3:[0-3]</i>         | <i>et-x/2/3</i>               | <i>et-x/2/3</i>                |
| 3                        | <i>xe-x/3/0:[0-3]</i>         | <i>et-x/3/0</i>               | <i>et-x/3/0</i>                |
|                          | <i>xe-x/3/1:[0-3]</i>         | <i>et-x/3/1</i>               | <i>et-x/3/1</i>                |
|                          | <i>xe-x/3/2:[0-3]</i>         | <i>et-x/3/2</i>               | <i>et-x/3/2</i>                |

Table 89: Interface Naming Convention for PTX10K-LC2101 MPC (Continued)

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
|                          | <i>xe-x/3/3:[0-3]</i>         | <i>et-x/3/3</i>               | <i>et-x/3/3</i>                |
| 4                        | <i>xe-x/4/0:[0-3]</i>         | <i>et-x/4/0</i>               | <i>et-x/4/0</i>                |
|                          | <i>xe-x/4/1:[0-3]</i>         | <i>et-x/4/1</i>               | <i>et-x/4/1</i>                |
|                          | <i>xe-x/4/2:[0-3]</i>         | <i>et-x/4/2</i>               | <i>et-x/4/2</i>                |
|                          | <i>xe-x/4/3:[0-3]</i>         | <i>et-x/4/3</i>               | <i>et-x/4/3</i>                |
| 5                        | <i>xe-x/5/0:[0-3]</i>         | <i>et-x/5/0</i>               | <i>et-x/5/0</i>                |
|                          | <i>xe-x/5/1:[0-3]</i>         | <i>et-x/5/1</i>               | <i>et-x/5/1</i>                |
|                          | <i>xe-x/5/2:[0-3]</i>         | <i>et-x/5/2</i>               | <i>et-x/5/2</i>                |
|                          | <i>xe-x/5/3:[0-3]</i>         | <i>et-x/5/3</i>               | <i>et-x/5/3</i>                |

**SEE ALSO**

[Line card \(MX10K-LC2101\)](#)

**Interface Naming Conventions for MIC-MACSEC-20GE**

By default, MIC-MACSEC-20GE operates in 1-Gigabit Ethernet mode. In this mode, 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. In this mode, the 10G ports physically maps to the front panel port 8 and 9 on the second PIC of the MIC (that is marked on the front panel of the MIC).

**NOTE:** In the 10G mode, only the marked ports are operational and other physical ports are disabled.

**Table 90: Interface Naming Convention for MIC-MACSEC-20GE**

| PIC   | 1-Gigabit Ethernet Interface | 10-Gigabit Ethernet Interface |
|-------|------------------------------|-------------------------------|
| PIC 0 | <i>ge-x/0/[0-9]</i>          | <i>xe-x/0/0</i>               |
| PIC 1 | <i>ge-x/1/[0-9]</i>          | <i>xe-x/1/0</i>               |
| PIC 2 | <i>ge-x/2/[0-9]</i>          | <i>xe-x/2/0</i>               |
| PIC 3 | <i>ge-x/3/[0-9]</i>          | <i>xe-x/3/0</i>               |

You should use the **pic-mode 10G** configuration command to set the PIC to operate in 10G mode. Both the PICs on a MIC must be configured in the same pic-mode, otherwise the configuration does not take effect. A chassis alarm is raised indicating a mis-configuration. Any mis-configuration will cause the PICs to assume default pic-mode, that is, to be in 20x1GE where all ports are in 1GE port speed.

**NOTE:** The 10-Gbps-capable ports (ports 8 and 9) of the 2x10GE/20x1GE MIC-MACSEC-20GE may show the link status as up while the peer side is down. In this case, it is recommended to disable auto-negotiation and set the speed to 1-Gbps on the peer side to bring the link up on the peer side.

The MIC-MACSEC-20GE MIC also provides 128-bit and 256-bit MACsec encryption on all the twenty 1GE and on the two 10GE ports in the following hardware configuration:

- Installed directly on the MX80 and MX104 routers
- Installed on MPC1, MPC2, MPC3, MPC2E, MPC3E, MPC2E-NG, and MPC3E-NG line cards on the MX240, MX480, and MX960 routers

By default, 128-bit MACsec encryption is supported.

The twenty 1-Gigabit Ethernet SFP ports distributes the ports across PIC0 and PIC1, that are logical PICs on the physical MIC. The two 10-Gigabit Ethernet SFP+ ports are physically located on PIC1. But, the 10-Gigabit interfaces are created by distributing the ports in either of the PICs.

**NOTE:**

- When the pic-mode is changed from 1-Gbps to 10-Gbps or vice versa, the Flexible PIC Concentrator (FPC) in MX240, MX480, MX960 routers and the Forwarding Engine Board (FEB) in MX80, MX104 routers undergoes an automatic bounce or a reboot.
- When the MIC-MACSEC-20GE is operating in the 10-Gbps mode, all the other 1-Gbps ports are disabled.

**SEE ALSO**

*Configuring Media Access Control Security (MACsec) on Routers*

*cipher-suite*

**Interface Naming Conventions for PTX10K-LC1201**

PTX10K-LC1201 is a fixed-configuration line card and contains 36 built-in ports. All ports support multiple port speeds of 400 Gbps, 200 Gbps, 100 Gbps, 50 Gbps, 40 Gbps, 25 Gbps and 10 Gbps and can be configured.

All the supported interfaces configured on the PTX10K-LC1201 line card follow the naming convention *et-fpc-slot/pic-slot/port-number.channel* where:

- Valid range for the FPC is 0 through 7.
- Valid range for the PIC is 0
- Valid range for the port is 0 through 35.
- Valid range for the channel is 0 through 7.

For example, the interface name *et-0/0/2* indicates any Gigabit Ethernet interface configured on port 2 of the PTX10K-LC1201 line card that is installed in the FPC slot 0.

[Table 60](#) lists the naming conventions used for interfaces on PTX10K-LC1201 for PTX10008 and PTX10016 routers. PTX10008 routers support 8 PTX10K-LC1201 line cards. PTX10016 routers support 16 PTX10K-LC1201 line cards.

Table 91: Interface Naming Convention for PTX10K-LC1201 line card

| 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 | <i>et-x/0/0:<br/>[0-7]</i>    | <i>et-x/0/0:<br/>[0-7]</i>    | <i>et-x/0/0</i>               | <i>et-x/0/0:<br/>[0-8]</i>    | <i>et-x/0/0<br/>et-x/0/0:[0-1]<br/>et-x/0/0:[0-3]</i> | <i>et-x/0/0:<br/>[0-2]</i>     | <i>et-x/0/0</i>                |
|       | <i>et-x/0/1:<br/>[0-7]</i>    | <i>et-x/0/1:<br/>[0-7]</i>    | <i>et-x/0/1</i>               | <i>et-x/0/1:<br/>[0-8]</i>    | <i>et-x/0/1<br/>et-x/0/1:[0-1]<br/>et-x/0/1:[0-3]</i> | <i>et-x/0/1:<br/>[0-2]</i>     | <i>et-x/0/1</i>                |
|       | <i>et-x/0/2:<br/>[0-7]</i>    | <i>et-x/0/2:<br/>[0-7]</i>    | <i>et-x/0/2</i>               | <i>et-x/0/2:<br/>[0-8]</i>    | <i>et-x/0/2<br/>et-x/0/2:[0-1]<br/>et-x/0/2:[0-3]</i> | <i>et-x/0/2:<br/>[0-2]</i>     | <i>et-x/0/2</i>                |
|       | <i>et-x/0/3:<br/>[0-7]</i>    | <i>et-x/0/3:<br/>[0-7]</i>    | <i>et-x/0/3</i>               | <i>et-x/0/3:<br/>[0-8]</i>    | <i>et-x/0/3<br/>et-x/0/3:[0-1]<br/>et-x/0/3:[0-3]</i> | <i>et-x/0/3:<br/>[0-2]</i>     | <i>et-x/0/3</i>                |
|       | <i>et-x/0/4:<br/>[0-7]</i>    | <i>et-x/0/4:<br/>[0-7]</i>    | <i>et-x/0/4</i>               | <i>et-x/0/4:<br/>[0-8]</i>    | <i>et-x/0/4<br/>et-x/0/4:[0-1]<br/>et-x/0/4:[0-3]</i> | <i>et-x/0/4:<br/>[0-2]</i>     | <i>et-x/0/4</i>                |
|       | <i>et-x/0/5:<br/>[0-7]</i>    | <i>et-x/0/5:<br/>[0-7]</i>    | <i>et-x/0/5</i>               | <i>et-x/0/5:<br/>[0-8]</i>    | <i>et-x/0/5<br/>et-x/0/5:[0-1]<br/>et-x/0/5:[0-3]</i> | <i>et-x/0/5:<br/>[0-2]</i>     | <i>et-x/0/5</i>                |

Table 91: Interface Naming Convention for PTX10K-LC1201 line card (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/0/6:<br>[0-7]            | et-x/0/6:<br>[0-7]            | et-x/0/6                      | et-x/0/6:<br>[0-8]            | et-x/0/6<br>et-x/0/6:[0-1]<br>et-x/0/6:[0-3]            | et-x/0/6:<br>[0-2]             | et-x/0/6                       |
|     | et-x/0/7:<br>[0-7]            | et-x/0/7:<br>[0-7]            | et-x/0/7                      | et-x/0/7:<br>[0-8]            | et-x/0/7<br>et-x/0/7:[0-1]<br>et-x/0/7:[0-3]            | et-x/0/7:<br>[0-2]             | et-x/0/7                       |
|     | et-x/0/8:<br>[0-7]            | et-x/0/8:<br>[0-7]            | et-x/0/8                      | et-x/0/8:<br>[0-8]            | et-x/0/8<br>et-x/0/8:[0-1]<br>et-x/0/8:[0-3]            | et-x/0/8:<br>[0-2]             | et-x/0/8                       |
|     | et-x/0/9:<br>[0-7]            | et-x/0/9:<br>[0-7]            | et-x/0/9                      | et-x/0/9:<br>[0-8]            | et-x/0/9<br>et-x/0/9:[0-1]<br>et-x/0/9:[0-3]            | et-x/0/9:<br>[0-2]             | et-x/0/9                       |
|     | et-x/0/10:<br>[0-7]           | et-x/0/10:<br>[0-7]           | et-x/0/10                     | et-x/0/10:<br>[0-8]           | et-x/0/10<br>et-x/0/10:<br>[0-1]<br>et-x/0/10:<br>[0-3] | et-x/0/10:<br>[0-2]            | et-x/0/10                      |

Table 91: Interface Naming Convention for PTX10K-LC1201 line card (Continued)

| PIC | 10-Gigabit Ethernet Interface | 25-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 50-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface                          | 200-Gigabit Ethernet Interface | 400-Gigabit Ethernet Interface |
|-----|-------------------------------|-------------------------------|-------------------------------|-------------------------------|---------------------------------------------------------|--------------------------------|--------------------------------|
|     | et-x/0/11:<br>[0-7]           | et-x/0/11:<br>[0-7]           | et-x/0/11                     | et-x/0/11:<br>[0-8]           | et-x/0/11<br>et-x/0/11:<br>[0-1]<br>et-x/0/11:<br>[0-3] | et-x/0/11:<br>[0-2]            | et-x/0/11                      |
|     | et-x/0/12:<br>[0-7]           | et-x/0/12:<br>[0-7]           | et-x/0/12                     | et-x/0/12:<br>[0-8]           | et-x/0/12<br>et-x/0/12:<br>[0-1]<br>et-x/0/12:<br>[0-3] | et-x/0/12:<br>[0-2]            | et-x/0/12                      |
|     | et-x/0/13:<br>[0-7]           | et-x/0/13:<br>[0-7]           | et-x/0/13                     | et-x/0/13:<br>[0-8]           | et-x/0/13<br>et-x/0/13:<br>[0-1]<br>et-x/0/13:<br>[0-3] | et-x/0/13:<br>[0-2]            | et-x/0/13                      |
|     | et-x/0/14:<br>[0-7]           | et-x/0/14:<br>[0-7]           | et-x/0/14                     | et-x/0/14:<br>[0-8]           | et-x/0/14<br>et-x/0/14:<br>[0-1]<br>et-x/0/14:<br>[0-3] | et-x/0/14:<br>[0-2]            | et-x/0/14                      |

Table 91: Interface Naming Convention for PTX10K-LC1201 line card (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/0/15:<br>[0-7]           | et-x/0/15:<br>[0-7]           | et-x/0/15                     | et-x/0/15:<br>[0-8]           | et-x/0/15<br>et-x/0/15:<br>[0-1]<br>et-x/0/15:<br>[0-3] | et-x/0/15:<br>[0-2]            | et-x/0/15                      |
|     | et-x/0/16:<br>[0-7]           | et-x/0/16:<br>[0-7]           | et-x/0/16                     | et-x/0/16:<br>[0-8]           | et-x/0/16<br>et-x/0/16:<br>[0-1]<br>et-x/0/16:<br>[0-3] | et-x/0/16:<br>[0-2]            | et-x/0/16                      |
|     | et-x/0/17:<br>[0-7]           | et-x/0/17:<br>[0-7]           | et-x/0/17                     | et-x/0/17:<br>[0-8]           | et-x/0/17<br>et-x/0/17:<br>[0-1]<br>et-x/0/17:<br>[0-3] | et-x/0/17:<br>[0-2]            | et-x/0/17                      |
|     | et-x/0/18:<br>[0-7]           | et-x/0/18:<br>[0-7]           | et-x/0/18                     | et-x/0/18:<br>[0-8]           | et-x/0/18<br>et-x/0/18:<br>[0-1]<br>et-x/0/18:<br>[0-3] | et-x/0/18:<br>[0-2]            | et-x/0/18                      |

Table 91: Interface Naming Convention for PTX10K-LC1201 line card (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/0/19:<br>[0-7]           | et-x/0/19:<br>[0-7]           | et-x/0/19                     | et-x/0/19:<br>[0-8]           | et-x/0/19<br>et-x/0/19:<br>[0-1]<br>et-x/0/19:<br>[0-3] | et-x/0/19:<br>[0-2]            | et-x/0/19                      |
|     | et-x/0/20:<br>[0-7]           | et-x/0/20:<br>[0-7]           | et-x/0/20                     | et-x/0/20:<br>[0-8]           | et-x/0/20<br>et-x/0/20:<br>[0-1]<br>et-x/0/20:<br>[0-3] | et-x/0/20:<br>[0-2]            | et-x/0/20                      |
|     | et-x/0/21:<br>[0-7]           | et-x/0/21:<br>[0-7]           | et-x/0/21                     | et-x/0/21:<br>[0-8]           | et-x/0/21<br>et-x/0/21:<br>[0-1]<br>et-x/0/21:<br>[0-3] | et-x/0/21:<br>[0-2]            | et-x/0/21                      |
|     | et-x/0/22:<br>[0-7]           | et-x/0/22:<br>[0-7]           | et-x/0/22                     | et-x/0/22:<br>[0-8]           | et-x/0/22<br>et-x/0/22:<br>[0-1]<br>et-x/0/22:<br>[0-3] | et-x/0/22:<br>[0-2]            | et-x/0/22                      |

Table 91: Interface Naming Convention for PTX10K-LC1201 line card (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:<br>[0-7]           | et-x/0/23:<br>[0-7]           | et-x/0/23                     | et-x/0/23:<br>[0-8]           | et-x/0/23<br>et-x/0/23:<br>[0-1]<br>et-x/0/23:<br>[0-3] | et-x/0/23:<br>[0-2]            | et-x/0/23                      |
|     | et-x/0/24:<br>[0-7]           | et-x/0/24:<br>[0-7]           | et-x/0/24                     | et-x/0/24:<br>[0-8]           | et-x/0/24<br>et-x/0/24:<br>[0-1]<br>et-x/0/24:<br>[0-3] | et-x/0/24:<br>[0-2]            | et-x/0/24                      |
|     | et-x/0/25:<br>[0-7]           | et-x/0/25:<br>[0-7]           | et-x/0/25                     | et-x/0/25:<br>[0-8]           | et-x/0/25<br>et-x/0/25:<br>[0-1]<br>et-x/0/25:<br>[0-3] | et-x/0/25:<br>[0-2]            | et-x/0/25                      |
|     | et-x/0/26:<br>[0-7]           | et-x/0/26:<br>[0-7]           | et-x/0/26                     | et-x/0/26:<br>[0-8]           | et-x/0/26<br>et-x/0/26:<br>[0-1]<br>et-x/0/26:<br>[0-3] | et-x/0/26:<br>[0-2]            | et-x/0/26                      |

Table 91: Interface Naming Convention for PTX10K-LC1201 line card (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/0/27:<br>[0-7]           | et-x/0/27:<br>[0-7]           | et-x/0/27                     | et-x/0/27:<br>[0-8]           | et-x/0/27<br>et-x/0/27:<br>[0-1]<br>et-x/0/27:<br>[0-3] | et-x/0/27:<br>[0-2]            | et-x/0/27                      |
|     | et-x/0/28:<br>[0-7]           | et-x/0/28:<br>[0-7]           | et-x/0/28                     | et-x/0/28:<br>[0-8]           | et-x/0/28<br>et-x/0/28:<br>[0-1]<br>et-x/0/28:<br>[0-3] | et-x/0/28:<br>[0-2]            | et-x/0/28                      |
|     | et-x/0/29:<br>[0-7]           | et-x/0/29:<br>[0-7]           | et-x/0/29                     | et-x/0/29:<br>[0-8]           | et-x/0/29<br>et-x/0/29:<br>[0-1]<br>et-x/0/29:<br>[0-3] | et-x/0/29:<br>[0-2]            | et-x/0/29                      |
|     | et-x/0/30:<br>[0-7]           | et-x/0/30:<br>[0-7]           | et-x/0/30                     | et-x/0/30:<br>[0-8]           | et-x/0/30<br>et-x/0/30:<br>[0-1]<br>et-x/0/30:<br>[0-3] | et-x/0/30:<br>[0-2]            | et-x/0/30                      |

Table 91: Interface Naming Convention for PTX10K-LC1201 line card (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/0/31:<br>[0-7]           | et-x/0/31:<br>[0-7]           | et-x/0/31                     | et-x/0/31:<br>[0-8]           | et-x/0/31<br>et-x/0/31:<br>[0-1]<br>et-x/0/31:<br>[0-3] | et-x/0/31:<br>[0-2]            | et-x/0/31                      |
|     | et-x/0/32:<br>[0-7]           | et-x/0/32:<br>[0-7]           | et-x/0/32                     | et-x/0/32:<br>[0-8]           | et-x/0/32<br>et-x/0/32:<br>[0-1]<br>et-x/0/32:<br>[0-3] | et-x/0/32:<br>[0-2]            | et-x/0/32                      |
|     | et-x/0/33:<br>[0-7]           | et-x/0/33:<br>[0-7]           | et-x/0/33                     | et-x/0/33:<br>[0-8]           | et-x/0/33<br>et-x/0/33:<br>[0-1]<br>et-x/0/33:<br>[0-3] | et-x/0/33:<br>[0-2]            | et-x/0/33                      |
|     | et-x/0/34:<br>[0-7]           | et-x/0/34:<br>[0-7]           | et-x/0/34                     | et-x/0/34:<br>[0-8]           | et-x/0/34<br>et-x/0/34:<br>[0-1]<br>et-x/0/34:<br>[0-3] | et-x/0/34:<br>[0-2]            | et-x/0/34                      |

**Table 91: Interface Naming Convention for PTX10K-LC1201 line card (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 |
|-----|-------------------------------|-------------------------------|-------------------------------|-------------------------------|------------------------------------------------------------------------|--------------------------------|--------------------------------|
|     | <i>et-x/0/35: [0-7]</i>       | <i>et-x/0/35: [0-7]</i>       | <i>et-x/0/35</i>              | <i>et-x/0/35: [0-8]</i>       | <i>et-x/0/35</i><br><i>et-x/0/35: [0-1]</i><br><i>et-x/0/35: [0-3]</i> | <i>et-x/0/35: [0-2]</i>        | <i>et-x/0/35</i>               |

**Interface Mapping and Modulation Format for ACX5448-D**

ACX5448-D routers support two CFP2-DCO optical modules (transceivers). For each optical module, one optical transport (*ot-*) interface is created. Thus, two *ot-* interfaces are created on this router. The ACX5448-D supports 100-Gigabit Ethernet (*et-*) interfaces. Two *et-* interfaces can be mapped to one *ot-* interface, depending on the rate (100 Gbps or 200 Gbps) that you configured for the CFP2 ports.

The optical interface to Ethernet interface mapping is displayed in the following table:

| "ot-" interface | Port Number | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------|-------------------|--------------------------|
| ot-0/2/0        | Port 38     | QPSK-100G         | et-0/2/0                 |
|                 |             | 8QAM-200G         | et-0/2/0<br>et-0/2/1     |
|                 |             | 16QAM-200G        | et-0/2/0<br>et-0/2/1     |
| ot-0/2/1        | Port 39     | QPSK-100G         | et-0/2/2                 |

*(Continued)*

| “ot-“ interface | Port Number | Modulation Format | Mapped “et” interface(s) |
|-----------------|-------------|-------------------|--------------------------|
|                 |             | 8QAM-200G         | et-0/2/2<br>et-0/2/3     |
|                 |             | 16QAM-200G        | et-0/2/2<br>et-0/2/3     |

The port panel of the ACX5448-D presents 36 SFP+ ports (0-35), two QSFP28 ports (36 and 37), and two CFP2-DCO ports (38 and 39). The port to logical PIC mapping is as follows:

- 0 to 35 ports represent PIC 0.
- 36 and 37 ports represent PIC 1.
- 38 and 39 ports represent PIC 2.

The QSFP28 port 36 (interface **et-0/1/0**) and the CFP2 port 38 (interface **et-0/2/0**) operate as multiplexer (also called mux) ports. Depending on the port speeds that you configure, the following behavior is observed:

- If you configure 8QAM or 16QAM modulation (200 Gbps) on **ot-0/2/0** (port 38), the **et-0/1/0** interface (port 36) is deleted.
- If you configure QPSK modulation (100 Gbps) on **ot-0/2/0** (port 38), then **et-0/2/1** on port 38 is disabled and **et-0/1/0** on port 36 is enabled. This means you can operate the multiplexed ports 36 (QSFP28) and 38 (CFP2) at 100-Gbps speeds.

When you start up the router, the two Ethernet interfaces on port 36 are disabled by default. However, the interface **et-0/2/0** (on port 38) is always available. You can enable the **et-0/1/0** interface (on port 36) by running the **set chassis fpc 0 cfp-to-et** command and restarting the chassis-control (You can restart the chassis control by using the **restart chassis-control** command). (This configuration deletes the interface **et-0/2/1** on port 38.) You can then multiplex interfaces **et-0/1/0** (port 36) and **et-0/2/0** (port 38) for a 200-Gbps operation. For more information on mux ports refer to [Port Panel of an ACX5448-D Router](#).

## SEE ALSO

| [cfp-to-et](#) | 707

## RELATED DOCUMENTATION

| [Configuring Port Speed on Routing Devices and Line Cards](#)

### Configuring Port Speed on Routing Devices and Line Cards

#### IN THIS SECTION

- [Configuring Port Speed | 418](#)
- [Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers | 419](#)
- [Configuring Port Speed on MIC-MRATE to Enable Different Port Speeds | 423](#)
- [Configuring Port Speed on MPC7E \(Multi-Rate\) to Enable Different Port Speeds | 427](#)
- [Configuring Port Speed on MX10003 MPC to Enable Different Port Speeds | 432](#)
- [Configuring Port Speed on MX204 to Enable Different Port Speeds | 436](#)
- [Configuring Port Speed on PTX10K-2101 MPC to Enable Different Port Speeds | 440](#)
- [Configure Port Speed on ACX5448-D and ACX5448-M Routers | 444](#)
- [Channelize Interfaces on ACX5448-D and ACX5448-M Routers | 445](#)
- [Configuring Port Speed on PTX10K-1201 line card to Enable Different Port Speeds | 448](#)

Use this topic for information about how to configure rate selectability on specific line cards. You can configure the speed of the port at the port level or at the PIC or MIC level.

#### Configuring Port Speed

Starting with Junos OS Release 15.1, some PICs support multiple port speeds. This procedure describes how to configure the port speed for these types of PICs.

To configure a PIC's port speed:

1. Navigate to the **[edit chassis]** hierarchy level.
2. Enter the **port-speed** statement at the **[edit chassis fpc slot-number pic pic-number port port-number]** hierarchy level.

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

- Specify the port speed that needs 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 port speed 10G;
user@host# set fpc fpc-slot pic pic-number port port-number port speed 40G;
user@host# set fpc fpc-slot pic pic-number port port-number port speed 100G;
```

## SEE ALSO

| *speed*

### Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers

PTX10003 routers (PTX10003-80C and PTX10003-160C) does not contain any pluggable PICs or TICs. You can directly plug-in the optics to the FPCs. Based on the optics, the interfaces are created with the respective interface naming conventions. The 40-Gigabit Ethernet, 100-Gigabit Ethernet, and 400-Gigabit Ethernet interfaces configured follow the naming convention **et-fpc-slot/pic-slot/port-number**. The 10-Gigabit Ethernet interfaces follow the naming convention **et-fpc-slot/pic-slot/port-number: [logical-port-number]**.

| Optic Device           | Interface speed | Interface Naming Format |
|------------------------|-----------------|-------------------------|
| QSFP56-DD-400GBASE-LR8 | 1x400G          | et-x/y/z                |
|                        | 4x100G          | et-x/y/z:0              |
|                        |                 | et-x/y/z:1              |
|                        |                 | et-x/y/z:2              |
|                        |                 | et-x/y/z:3              |

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 *Unused*

for more details). When using port 4 as 400G, port 3 has to be configured with total bandwidth of less than 100G 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 than 100-Gbps, and the middle port (2 between 0~4 or 7 between 5~9) must be set to unused. For example, you can set et-0/0/0 to 400G, et-0/0/1 to 100G or less, but et-0/0/2 must be set to unused.

To view the port panel information, refer to [PTX10003 Port Panel](#).

For each PIC, maximum speed supported on the respective port is limited to:

| Port Number | Speed  |
|-------------|--------|
| Port 0      | 400G   |
| Port 1      | 100G   |
| Port 2      | Unused |
| Port 3      | 100G   |
| Port 4      | 400G   |
| Port 5      | 400G   |
| Port 6      | 100G   |
| Port 7      | Unused |
| Port 8      | 100G   |
| Port 9      | 400G   |

This topic describes the guidelines to be considered before you configure 4x100 Gbps on PTX10003 routers:

- When you configure port 0 with 4x100 Gbps speed, you must configure port 1 and port 2 as unused.
- When you configure port 4 with 4x100 Gbps speed, you must configure port 2 and port 3 as unused.
- When you configure port 5 with 4x100 Gbps speed, you must configure port 6 and port 7 as unused.
- When you configure port 9 with 4x100 Gbps speed, you must configure port 7 and port 8 as unused.

Table 61 lists the guidelines to configure 4x100 Gbs on the PTX10003 routers in a tabular format.

**Table 92: Configuration Guidelines to configure 1x400 Gbps on PTX10003 routers**

| Ports with speed 4x100 gbps | Unusued Ports     |
|-----------------------------|-------------------|
| Port 0                      | Port 1 and Port 2 |
| Port 4                      | Port 2 and Port 3 |
| Port 5                      | Port 6 and Port 7 |
| Port 9                      | Port 7 and Port 8 |

By default, the PIC comes up with default interfaces in 100-Gigabit Ethernet mode. To configure 400-Gbps speed on PTX10003 routers (PTX10003-80C and PTX10003-160C):

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

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

For example:

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

2. To configure 400-Gbps speed at the specific port, configure the **speed** statement for the desired ports.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed 400G
```

**NOTE:** You can only configure 400-Gbps speed on ports 0, 4, 5, 9 within each logical PIC.

For example:

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port 0 speed 400G
```

3. Since port 0 is configured to 400-Gbps speed, you can only configure et-0/0/1 to 100G or less, and et-0/0/2 to unused in the same logical PIC.

For example:

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port 1 speed 100G
user@host# set port 2 unused
```

4. Commit your configuration changes.

On successful commit, et-1/1/0 is created with 400-Gbps, et-1/1/1 is created with 100-Gbps speed, and no interfaces will be created on port 2.

## SEE ALSO

[speed \(Ethernet\) | 1137](#)

[Unused | 1220](#)

[number-of-sub-ports | 1042](#)

## Configuring Port Speed on MIC-MRATE to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Port Speed on MIC-MRATE at MIC Level | 423](#)
- [Configuring Port Speed on MIC-MRATE at Port Level | 425](#)

Rate selectability enables you to configure the port speed either at the port level or at the MIC level. To configure all ports to operate at the same speed, you configure rate selectability at the MIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the MIC level, use the **pic-mode** statement and specify the port speed. To configure different port speeds for each port, you configure rate selectability at the port level, in which case only the ports that are configured are enabled. To configure rate selectability at the port level, use the **speed** statement to specify the speed of individual ports. This topic describes how to configure port speeds at the port level and at the MIC level.

**NOTE:** The **pic** in the configuration commands refers to the MRATE MIC. To specify *pic-number*, specify the MRATE MIC slot. For instance, when you use the **edit chassis fpc *fpc-slot-number* pic *pic-slot-number*** statement, specify the MPC slot number and the MIC-MRATE slot number.

### *Configuring Port Speed on MIC-MRATE at MIC Level*

To configure all ports to operate at the same speed, you configure rate selectability at the MIC level. The default port speed is 10 Gbps for all ports. When you configure rate selectability at the MIC level, all the ports of the MIC that support the configured speed operate at that speed. To prevent oversubscription and ensure a guaranteed bandwidth, you can specify the number of active ports that operate at the configured speed by using the **number-of-ports *number-of-active-physical-ports*** statement. MIC-MRATE supports port speeds of 10 Gbps, 40 Gbps, and 100 Gbps.

For MPC8E, you can only configure 4 ports of the 12 MIC-MRATE ports with 100 Gbps port speed and the other ports are disabled. So, if you configure **100G** as the operating speed for ports 0, 1, 6, and 7, then the other ports are disabled on MPC8E. Similarly, when you configure the port speed as 100 Gbps at the MIC level on MPC9E, you can only configure 8 ports of the 12 MIC-MRATE ports to operate with that speed. So, if you configure **100G** as the operating speed for ports 0, 1, 2, 3, 6, 7, 8, and 9, then the other ports are disabled on MPC9E. However, enabling port speed of 40 Gbps or 10 Gbps at the MIC level, enables all ports and sets the desired port speed on all ports.

To configure rate selectability at the MIC level:

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

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

2. Configure the `pic-mode pic-speed` statement to set the operating speed for the MIC. All ports of the MIC that support the configured speed operate at the configured speed. Values for the `pic-speed` option include: **10G**, **40G**, and **100G**.

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

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set pic-mode 10G
```

3. (Optional) To prevent oversubscription, you can choose to configure the number of active ports that operate at the port speed configured in Step 2. For information about the number of active ports and specific port numbers on MPC7E-MRATE, MPC8E, and MPC9E see [Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription](#).

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set number-of-ports number-of-active-physical-ports
```

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set number-of-ports 8
```

4. Verify the configuration.

```
[edit chassis fpc 4 pic 0]
user@host# show
pic-mode 10G;
number-of-ports 8;
```

5. Commit your configuration changes.

In this example, you have configured 8 ports on MIC-MRATE with port speed of 10 Gbps. The other ports are disabled.

### ***Configuring Port Speed on MIC-MRATE at Port Level***

To configure different port speeds for each port, you configure rate selectability at the port level. Only the ports that are configured are enabled. Other ports are automatically disabled. Configuring rate selectability at the port level provides you the flexibility of operating individual ports of the MIC at different supported speeds. For example, on MPC9E with MIC-MRATE, you can configure four 100-Gigabit Ethernet interfaces on ports 0, 1, 6, and 7 and two 40-Gigabit Ethernet interfaces on ports 3 and 8. You can use breakout transceivers to configure each 40-Gigabit Ethernet interfaces as four 10-Gigabit Ethernet interfaces.

**NOTE:** When you change the port speed at the port level, you must reset the MIC for the configuration to take effect. To reset the MIC, use the **request chassis mic mic-slot *mic-slot-number* fpc-slot *fpc-slot-number* (offline | online)** command to reset the MIC and apply your configuration changes. Alternatively, you can also restart the MPCs. However, MPC restart takes longer as it affects all the PFEs. An alarm is generated indicating the change in port speed. For guidelines on configuring rate selectability, see *Guidelines for Configuring Port Speed on Routing Devices*.

To configure rate selectability at the port level:

1. In configuration mode, navigate to the **[edit chassis fpc *fpc-slot* pic *pic-number*]** hierarchy level.

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

For example:

```
[edit]
user@host# edit chassis fpc 4 pic 0
```

2. To indicate the speed at which the ports operate, configure the **speed** statement for the desired ports. According to your requirements, you can choose the **10G**, **40G**, or **100G** speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G | 40G | 100G)
```

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set port 0 speed 100G
user@host# set port 1 speed 100G
user@host# set port 3 speed 40G
user@host# set port 6 speed 100G
user@host# set port 7 speed 100G
user@host# set port 8 speed 40g
```

**NOTE:** All the twelve ports of MIC-MRATE support 10-Gbps and 40-Gbps port speeds. On MPC8E with MIC-MRATE, you can configure 4 ports out of the twelve MIC-MRATE ports with a port speed of 100 Gbps. On MPC9E with MIC-MRATE, you can configure 8 ports out of the twelve MIC-MRATE ports with a port speed of 100 Gbps.

### 3. Verify the configuration.

```
[edit chassis fpc 4 pic 0]
user@host# show
port 0 {
 speed 100g;
}
port 1 {
 speed 100g;
}
port 3 {
 speed 40g;
}
port 6 {
 speed 100g;
}
port 7 {
 speed 100g;
}
port 8 {
 speed 40g;
}
```

### 4. Commit your configuration changes.

In this example, you have configured 4 ports on MIC-MRATE with port speed of 100 Gbps and 2 ports with port speed of 40 Gbps. The total capacity per MIC, based on this configuration, is 480 Gbps. MIC-MRATE has two Packet Forwarding Engines. The forwarding capacity for each Packet Forwarding Engine is 400 Gbps for MPC9E and 240 Gbps for MPC8E. The configured value does not exceed the forwarding capacity and so is a valid configuration.

## RELATED DOCUMENTATION

*number-of-ports*

*pic-mode*

*speed*

*Port Speed for Routing Devices*

## Configuring Port Speed on MPC7E (Multi-Rate) to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Port Speed at PIC Level | 428](#)
- [Configuring Port Speed at Port Level | 429](#)

Each of the six ports of PIC 0 and PIC 1 of an MPC7E-MRATE MPC supports port speeds of 10 Gbps and 40 Gbps. However, only ports 2 and 5 of both the PICs support port speed of 100 Gbps. Because the MPC7E-MRATE MPC is rate-selectable, you can choose to configure all supported ports of the MPC to operate at the same supported speed or configure all the ports at different supported speeds.

You configure rate selectability at the PIC level if you intend to operate all the ports of the MPC7E-MRATE MPC at the same speed. That is, you can choose to configure the PIC to operate at a supported speed, and then all the *supported ports* of the PIC operate at the configured speed. For example, if you choose to configure PIC 0 at 100-Gbps speed, only ports 2 and 5 of PIC 0 operate at 100-Gbps speed, while the other ports of the PIC are disabled. Similarly, if you choose to configure PIC 0 at 10-Gbps or 40-Gbps speed, all the ports of the PIC are enabled to operate at those speeds. Additionally, you can prevent oversubscription by specifying the number of active physical ports that operate at 10-Gbps, 40-Gbps, and 100-Gbps speeds.

You configure rate selectability at the port level if you intend to operate different ports of the MPC7E-MRATE MPC at different supported speeds. That is, you configure each port to operate at a supported speed.

**NOTE:** The MPC7E-MRATE MPC supports an aggregate bandwidth of 480 Gbps, and each of the two PICs supports a bandwidth limit of 240 Gbps. If the aggregate port capacity configured exceeds 240 Gbps per PIC, the configuration is not supported.

### *Configuring Port Speed at PIC Level*

To configure rate selectability at the PIC level:

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

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

For example:

```
[edit]
user@host# edit chassis fpc 4 pic 0
```

2. Configure the `pic-mode` statement to set the operating speed for the PIC's ports. According to your requirements, you can choose from the options **10G**, **40G**, or **100G**.

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

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set pic-mode 10G
```

3. (Optional) To prevent oversubscription, you can choose to configure the number of ports that operate at the mode configured in Step 2.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set number-of-ports number-of-active-physical-ports
```

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set number-of-ports 6
```

#### 4. Verify the configuration.

```
[edit chassis fpc 4 pic 0]
user@host# show
pic-mode 10G;
number-of-ports 6;
```

#### 5. Commit your configuration changes.

If the **number-of-ports** statement is *not* configured, all the ports that support the speed configured in Step "2" are enabled. That is, depending on that selection, ports 0 through 5 are enabled for speeds of 10-gigabit or 40-gigabit, while ports 2 and 5 are enabled for 100-gigabit. [Table 62](#) lists the physical ports that are enabled when the **number-of-ports** statement is configured.

**Table 93: Active Physical Ports on MPC7E-MRATE MPC Based on the number-of-ports Configuration**

| Ports Configured (number-of-ports Statement) | Active Physical Ports for Different Configured Speeds |                  |             |
|----------------------------------------------|-------------------------------------------------------|------------------|-------------|
|                                              | 10-Gigabit                                            | 40-Gigabit       | 100-Gigabit |
| 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        |

*Configuring Port Speed at Port Level*

To configure rate selectability at the port level:

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

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

For example:

```
[edit]
user@host# edit chassis fpc 4 pic 0
```

2. To indicate the speed at which the ports operate, configure the `speed` statement for the desired ports. According to your requirements, you can choose the **10g**, **40g**, or **100g** speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10g | 40g | 100g)
```

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set port 0 speed 10g
user@host# set port 1 speed 10g
user@host# set port 2 speed 100g
user@host# set port 3 speed 40g
```

**NOTE:** All the six ports of PIC 0 and PIC 1 of an MPC7E-MRATE MPC support 10-Gbps and 40-Gbps port speeds. However, only ports 2 and 5 of both the PICs support 100-Gbps speed.

3. Verify the configuration.

```
[edit chassis fpc 4 pic 0]
user@host# show
port 0 {
 speed 10g;
}
port 1 {
 speed 10g;
```

```

}
port 2 {
 speed 100g;
}
port 3 {
 speed 40g;
}

```

#### 4. Commit your configuration changes.

**NOTE:** Note the following when configuring rate selectability on an MPC7E-MRATE MPC:

- If rate selectability is not configured, all ports of the MPC7E-MRATE MPC operate as four 10-Gigabit Ethernet interfaces by default. Therefore, when booting the MPC:
  - If rate selectability is not configured or if invalid port speeds are configured, each port operates as four 10-Gigabit Ethernet interfaces. An alarm is generated to indicate that the ports of the MPC7E-MRATE MPC are operating as four 10-Gigabit Ethernet interfaces.
  - If valid port speeds are configured, the MPC PICs operate at the configured speed.
- When you change an existing port speed configuration at the port level, you must reset the MPC7E-MRATE PIC for the configuration to take effect. An alarm is generated indicating the change in port speed configuration.
- When you change an existing port speed configuration with an *invalid* port speed configuration, an alarm is generated indicating that the port speed configuration is invalid. The MPC continues to operate using the previously configured valid port speed configuration. However, if the MPC or PIC is restarted with the committed invalid port configuration, all ports of the MPC operate as four 10-Gigabit Ethernet interfaces by default.
- You cannot configure rate selectability at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations.
- When you configure rate selectability at the port level, only the configured ports are enabled. Other ports are disabled.
- Logical interfaces can be created only on ports that are enabled.

## RELATED DOCUMENTATION

*pic-mode*

*speed*

| *number-of-ports*

## Configuring Port Speed on MX10003 MPC to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Port Speed on MX10003 MPC at MIC/PIC Level | 432](#)
- [Configuring Port Speed on MX10003 MPC at Port Level | 434](#)

Rate selectability enables you to configure the port speed either at the port level or at the MIC level. To configure all ports to operate at the same speed, configure rate selectability at the MIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the MIC level, use the **pic-mode** statement and specify the port speed. To configure different port speeds for each port, configure rate selectability at the port level, in which case only the ports that are configured are enabled. To configure rate selectability at the port level, use the **speed** statement to specify the speed of individual ports. This topic describes how to configure port speeds at the port level and at the MIC or PIC level.

**NOTE:** Regardless of the line card— MIC (PIC1) or fixed-port PIC (PIC0) installed —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.

### *Configuring Port Speed on MX10003 MPC at MIC/PIC Level*

To configure all ports to operate at the same speed, configure rate selectability at the MIC or PIC level. When you configure rate selectability at the MIC or PIC level, all the ports of the MIC that support the configured speed operate at that speed. To prevent oversubscription and to ensure a guaranteed bandwidth, specify the number of active ports that operate at the configured speed by using the **number-of-ports** *number-of-active-physical-ports* statement. The MX10003 MPC supports port speeds of 10 Gbps, 40 Gbps, and 100 Gbps.

To configure rate selectability at the MIC/PIC level:

1. In configuration mode, navigate to the **[edit chassis fpc *fpc-slot* pic *pic-number*]** hierarchy level.

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

For example:

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

2. Configure the **pic-mode *pic-speed*** statement to set the operating speed for the MIC. All ports of the MIC that support the configured speed operate at the configured speed. Values for the ***pic-speed*** option are 10G, 40G, and 100G.

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

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set pic-mode 10G
```

3. (Optional) To prevent oversubscription, you can choose to configure the number of active ports that operate at the port speed configured in Step 2. For information about the number of active ports and specific port numbers on the MX10003 MPC , see [Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription on MX10003 MPC](#).

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set number-of-ports number-of-active-physical-ports
```

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set number-of-ports 8
```

4. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
pic-mode 10G;
number-of-ports 8;
```

5. Commit your configuration changes.

In this example, you have configured 8 ports on the MPC with port speed of 10 Gbps. The other ports are disabled.

### ***Configuring Port Speed on MX10003 MPC at Port Level***

To configure different port speeds for each port, you configure rate selectability at the port level. Only the ports that are configured are enabled. Other ports are automatically disabled. Configuring rate selectability at the port level provides you the flexibility of operating individual ports of the PIC at different supported speeds.

If you want to configure a port speed of 40 Gbps on the MIC and the fixed-port PIC, you can choose any of the following example configurations:

- Configure one port of the fixed-port PIC as a 40-Gigabit Ethernet interface and three ports of the MIC as 40-Gigabit Ethernet interfaces.
- Configure two ports of the fixed-port PIC as 40-Gigabit Ethernet interfaces and three ports of the MIC as 40-Gigabit Ethernet interfaces.
- Configure three ports of the MIC as 40-Gigabit Ethernet interfaces and two ports of the fixed-port PIC as 40-Gigabit Ethernet interfaces.
- Configure four ports of the MIC as 40-Gigabit Ethernet interfaces only.

**NOTE:** While configuring rate selectability, when you switch to PIC mode from port mode or vice-versa, the PIC is reset automatically. However, when you change the port speed at the port level, the PIC has to be reset by executing the **request chassis pic pic-slot mic-slot-number fpc-slot fpc-slot-number (offline | online)** command. For guidelines on configuring rate selectability, see *Guidelines for Configuring Port Speed on Routing Devices*.

To configure rate selectability at the port level:

1. In configuration mode, navigate to the **[edit chassis fpc fpc-slot pic pic-number]** hierarchy level.

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

For example:

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

2. To indicate the speed at which the ports operate, configure the **speed** statement for the desired ports. According to your requirements, you can choose the 10G, 40G, and 100G speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G | 40G | 100G)
```

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set port 0 speed 10G
user@host# set port 1 speed 10G
user@host# set port 3 speed 40G
```

3. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
port 0 {
 speed 10g;
}
port 1 {
 speed 10g;
}
port 3 {
 speed 40g;
}
```

4. Commit your configuration changes.

In this example, you have configured 2 ports with port speed of 10 Gbps and 1 port with port speed of 40 Gbps.

**NOTE:** Starting in Junos OS Release 18.1R1, the 10-Gbps port can operate in 1-Gbps mode also using the *speed (Gigabit Ethernet interface)* configuration statement at Gigabit Ethernet interface level. Refer to *MX10003 MPC Port Speed Overview* for more details.

## RELATED DOCUMENTATION

*number-of-ports*

*pic-mode*

*speed*

*Port Speed for Routing Devices*

## Configuring Port Speed on MX204 to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Port Speed on MX204 at PIC Level | 437](#)
- [Configuring Port Speed on MX204 at Port Level | 438](#)

Rate selectability enables you to configure the port speed either at the port level or at the PIC level. To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the *pic-mode* statement and specify the port speed. To configure different port speeds for each port, configure rate selectability at the port level, in which case only the ports that are configured are enabled. To configure rate selectability at the port level, use the *speed* statement to specify the speed of individual ports. This topic describes how to configure port speeds at the port level and at the PIC level.

Note the following caveats while configuring rate selectability on the MX204 routers:

- Regardless of the line card— MIC (PIC1) or fixed-port PIC (PIC0) installed —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.
- If rate selectability is not configured, all ports of the MX204 router operate as 10-Gigabit Ethernet interfaces.
- 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. For more information, see *MX204 Router Port Speed Overview*.
- The heterogeneous mode is supported only on port mode.
- When you configure rate selectability at the port level, only the configured ports are active. Other ports are disabled.

- 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 rate selectability at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations.

### ***Configuring Port Speed on MX204 at PIC Level***

To configure all ports to operate at the same speed, you configure rate selectability at the PIC level. When you configure rate selectability at the PIC level, all the ports of the PIC that support the configured speed operate at that speed. To prevent oversubscription and ensure a guaranteed bandwidth, you can specify the number of active ports that operate at the configured speed by using the **number-of-ports** *number-of-active-physical-ports* statement. The MX204 has four rate-selectable ports (referred to as PIC 0 ports) that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet port, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The MX204 also has eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports).

To configure rate selectability at the PIC level:

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

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

For example:

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

2. Configure the **pic-mode** *pic-speed* statement to set the operating speed for the PIC. All ports of the PIC that support the configured speed operate at the configured speed. Values for the *pic-speed* option are **10G**, **40G**, and **100G**.

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

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set pic-mode 10G
```

3. (Optional) To prevent oversubscription, you can choose to configure the number of active ports that operate at the port speed configured in Step 2. For information about the number of active ports and specific port numbers on the MX204 routers see [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router](#).

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set number-of-ports number-of-active-physical-ports
```

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set number-of-ports 4
```

4. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
pic-mode 10G;
number-of-ports 4;
```

5. Commit your configuration changes.

In this example, you have configured 4 ports on the PIC0 with port speed of 10 Gbps.

#### ***Configuring Port Speed on MX204 at Port Level***

To configure different port speeds for each port, you configure rate selectability at the port level. Only the ports that are configured are enabled. Other ports are automatically disabled. Configuring rate selectability at the port level provides you the flexibility of operating individual ports of the PIC at different supported speeds.

**NOTE:** When you change the port speed at the port level, you must reset the PIC for the configuration to take effect. Resetting the PIC takes several minutes and affects all the Packet Forwarding Engines. To avoid this, use the **request chassis pic *pic-slot mic-slot-number* fpc-slot *fpc-slot-number* (offline | online)** command to reset the PIC and apply your configuration changes. An alarm is generated indicating the change in port speed. For guidelines on configuring rate selectability, see *Guidelines for Configuring Port Speed on Routing Devices*.

To configure rate selectability at the port level:

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

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

For example:

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

2. To indicate the speed at which the ports operate, configure the `speed` statement for the desired ports. According to your requirements, you can choose the **10G**, **40G**, or **100G** speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G | 40G | 100G)
```

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set port 0 speed 100G
user@host# set port 1 speed 40G
user@host# set port 2 speed 40G
user@host# set port 3 speed 10G
```

3. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
port 0 {
 speed 100g;
}
port 1 {
 speed 40g;
}
port 2 {
 speed 40g;
}
port 3 {
 speed 10g;
```

```
}
}
```

#### 4. Commit your configuration changes.

In this example, you have configured 2 ports on the PIC0 with port speed of 40 Gbps, 1 port with port speed of 10 Gbps, and 1 port with port speed of 100 Gbps.

**NOTE:** Starting in Junos OS Release 18.1R1, the 10-Gbps port can operate in 1-Gbps mode also using the *speed (Gigabit Ethernet interface)* configuration statement at Gigabit Ethernet interface level. Refer to *MX10003 MPC Port Speed Overview* for more details.

## RELATED DOCUMENTATION

[number-of-ports](#)

[pic-mode](#)

[speed](#)

[Port Speed for Routing Devices](#)

[MX204 Router Overview](#)

[request chassis pic](#)

## Configuring Port Speed on PTX10K-2101 MPC to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Port Speed on PTX10K-2101 MPC at PIC Level | 441](#)
- [Configuring Port Speed on PTX10K-LC2101 MPC at Port Level | 442](#)

Rate selectability enables you to configure the port speed either at the port level or at the PIC level. To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the **pic-mode** statement and specify the port speed. To configure different port speeds for each port, configure rate selectability at the port level, in which case only the ports that are configured are

enabled. To configure rate selectability at the port level, use the **speed** statement to specify the speed of individual ports. This topic describes how to configure port speeds at the port level and at the PIC level.

### *Configuring Port Speed on PTX10K-2101 MPC at PIC Level*

To configure all ports to operate at the same speed, configure rate selectability at the PIC level. When you configure rate selectability at the PIC level, all the ports of the PIC that support the configured speed operate at that speed. To prevent oversubscription and to ensure a guaranteed bandwidth, specify the number of active ports that operate at the configured speed by using the **number-of-ports** *number-of-active-physical-ports* statement. The PTX10K-LC2101 MPC supports port speeds of 10 Gbps, 40 Gbps, and 100 Gbps.

To configure rate selectability at the PIC level:

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

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

For example:

```
[edit]
user@host# edit chassis fpc 5 pic 2
```

2. Configure the **pic-mode** *pic-speed* statement to set the operating speed for the PIC. All ports of the PIC that support the configured speed operate at the configured speed. Values for the *pic-speed* option are 10G, 40G, and 100G.

**NOTE:** When you configure the **pic-mode** as 100 Gbps and the Packet Forwarding Engine bandwidth is 240 Gbps, only the first two ports support 100 Gbps. The other ports are disabled.

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

For example:

```
[edit chassis fpc 5 pic 2]
user@host# set pic-mode 10G
```

3. (Optional) To prevent oversubscription, you can choose to configure the number of active ports that operate at the port speed configured in Step 2.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set number-of-ports number-of-active-physical-ports
```

For example:

```
[edit chassis fpc 5 pic 2]
user@host# set number-of-ports 2
```

4. Verify the configuration.

```
[edit chassis fpc 5 pic 2]
user@host# show
pic-mode 10G;
number-of-ports 2;
```

5. Commit your configuration changes.

In this example, you have configured 2 ports on the MPC with port speed of 10 Gbps. The other ports are disabled.

#### *Configuring Port Speed on PTX10K-LC2101 MPC at Port Level*

To configure different port speeds for each port, you configure rate selectability at the port level. Only the ports that are configured are enabled. Other ports are automatically disabled. Configuring rate selectability at the port level provides you the flexibility of operating individual ports of the PIC at different supported speeds.

**NOTE:** While configuring rate selectability, when you switch to PIC mode from port mode or vice-versa, the PIC is reset automatically. However, when you change the port speed at the port level, the PIC has to be reset by executing the **request chassis pic pic-slot *mic-slot-number* fpc-slot *fpc-slot-number* (offline | online)** command. For guidelines on configuring rate selectability for PTX10K-LC2101, see *Guidelines for Configuring Port Speed on Routing Devices*.

To configure rate selectability at the port level:

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

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

For example:

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

2. To indicate the speed at which the ports operate, configure the **speed** statement for the desired ports. According to your requirements, you can choose the 10G, 40G, and 100G speed options.

**NOTE:** If you configure the **speed** as 100 Gbps for 3 ports and the Packet Forwarding Engine bandwidth is 240 Gbps, an alarm is raised as it is an invalid configuration. The value of only the first two ports support 100 Gbps. The other ports are disabled.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G | 40G | 100G)
```

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set port 0 speed 10G
user@host# set port 1 speed 10G
user@host# set port 3 speed 40G
```

3. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
port 0 {
 speed 10g;
}
port 1 {
 speed 10g;
}
```

```
port 3 {
 speed 40g;
}
```

#### 4. Commit your configuration changes.

In this example, you have configured 2 ports with port speed of 10 Gbps and 1 port with port speed of 40 Gbps.

## RELATED DOCUMENTATION

---

*number-of-ports*

---

*pic-mode*

---

*speed*

---

*Port Speed for Routing Devices*

### Configure Port Speed on ACX5448-D and ACX5448-M Routers

To configure speeds on different ports, you configure rate selectability at the port level. Configuring rate selectability at the port level provides you the flexibility of operating individual ports of a PIC at different supported speeds. The ACX5448-D router contains 36 SFP+ ports (**0** through **35**), two 100-Gigabit Ethernet QSFP28 ports (**36** and **37**), and two CFP2-DCO ports (**38** and **39**). When you start up the ACX5448-D router, the two Ethernet interfaces on port 36 are disabled by default. The ACX5448-M router contains 44 SFP+ ports (port **0–43** on PIC 0) and 6 QSFP28 ports (ports **0–5** on PIC 1).

This topic describes how to configure speeds at the port level. The ACX5448-D supports port speeds of 10-Gbps, 25-Gbps, 40-Gbps, and 100-Gbps.

To configure rate selectability at the port level on ACX5448-D:

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

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

For example:

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

2. To indicate the speed at which the ports operate, configure the **speed** statement for the specified ports. We know that the ports on PIC 1 support speeds of 10 Gbps, 25 Gbps, 40 Gbps, and 100 Gbps. According to your requirement, you can choose any of the speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G | 25G | 40G | 100G)
```

For example:

```
[edit chassis fpc 0 pic 1]
user@host# set port 0 speed 10G
user@host# set port 1 speed 25G
```

3. Verify the configuration.

```
[edit chassis fpc 0 pic 1]
user@host# show
port 0 {
 speed 10g;
}
port 1 {
 speed 25g;
}
```

4. Commit your configuration changes.

In this example, you have configured 25-Gbps speed on one port and 10-Gbps speed on another on the ACX5448-D router.

## SEE ALSO

| [Introduction to Port Speed](#) | 0

### Channelize Interfaces on ACX5448-D and ACX5448-M Routers

The ACX5448 router has two SKUs, ACX5448-D and ACX5448-M. The ACX5448-D router has 40 network ports, categorized as 36 SFP+/SFP ports, 2 QSFP28 ports, and 2 CFP2-DCO ports. You can configure 36 ports (ports **0–35**) as 1-Gigabit or 10-Gigabit Ethernet interfaces. The QSFP28 ports (ports **36** and **37**) support 100-Gbps and 40-Gbps speeds; you can channelize these ports into four 25-Gigabit

or four 10-Gigabit Ethernet interfaces, respectively, using the **set chassis fpc fpc slot pic pic slot port port-number speed speed** command. Each of the CFP2-DCO ports (ports **38** and **39**) supports up to 200-Gbps speed.

In the Junos OS CLI, we have mapped the ports on the ACX5448-D to logical PICs in the following manner:

- Ports **0** through **35** (with the xe- interface type) represent PIC 0.
- Ports **36** and **37** (with the et- interface type) represent PIC 1.
- Ports **38** and **39** (with the ot- interface type) represent PIC 2.

On the ACX5448-M router, has a total of 50 network ports, categorized as 44 SFP+ and 6 QSFP28 ports. You can configure 44 ports (port **0–43** on PIC 0) as 1-Gigabit or 10-Gigabit Ethernet interfaces. The QSFP28 ports (ports **0–5** on PIC 1) support 100-Gbps and 40-Gbps speeds; you can channelize these ports into four 25-Gigabit Ethernet or four 10-Gigabit Ethernet interfaces, respectively, using the **set chassis fpc fpc slot pic pic slot port port-number speed speed** command. By default, each of the QSFP28 ports (ports **0–5**) supports 100-Gbps speed.

In the Junos OS CLI, we have mapped the ports on the ACX5448-M to logical PICs in the following manner:

- Ports **0** through **43** mapped to PIC 0 (interfaces xe-0/0/0 through xe-/0/0/43).
- Ports **44** through **49** mapped to PIC 1 (interfaces et-0/0/0 through et-/0/0/5).

All the channelized interfaces in a port use the format *fpc/pic/port:channel-number*—where *channel-number* can be a value from 0 through 3—and have the same port properties. By default, the port speed is 100-Gbps on PIC 1 (ports **36** and **37**).

When you start up the router, the et-0/1/0 interface on port **36** is not created by default. However, the interface et-0/2/0 (on port **38**) is always available. You can enable the et-0/1/0 interface (on port **36**) by configuring the **set chassis fpc 0 cfp-to-et** command and restarting the FPC by executing the **restart chassis-control** command. (This configuration deletes the interface et-0/2/1 on port **38**.)

As the default speed is 100-Gbps, we must first change the speed to 40-Gbps, and then channelize that port into four 10-Gbps interfaces. To channelize the ports, manually configure the port speed using the **set chassis fpc slot-number pic pic-number port port-number speed speed** command where the speed can be set to 4x10-Gbps or 4x25-Gbps.

To channelize an individual port:

1. To configure an individual 100-Gigabit Ethernet (et-) port to operate as four 25-Gbps interfaces, specify the port number and speed:

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

For example, to configure port **1** (et-0/1/1 interface) to operate as four 25-Gbps interfaces:

```
[edit]
user@host# set chassis fpc 0 pic 1 port 1 speed 25g
```

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

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

3. To revert the four 25-Gbps channelized interfaces to operate as a single default 100-Gbps interface, delete the speed statement:

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

For example, to return port **1** from the 25-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 1]
user@host# delete port 1 speed 25g
```

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

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

## SEE ALSO

*fpc*

*pic*

## Configuring Port Speed on PTX10K-1201 line card to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Port Speed on PTX10K-1201 line card at PIC Level | 448](#)
- [Configuring Port Speed on PTX10K-LC1201 line card at Port Level | 449](#)

Rate selectability enables you to configure the port speed either at the port level or at the PIC level. To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the **pic-mode** statement and specify the port speed. To configure different port speeds for each port, configure rate selectability at the port level, in which case only the ports that are configured are enabled. To configure rate selectability at the port level, use the **speed** statement to specify the speed of individual ports. This topic describes how to configure port speeds at the port level and at the PIC level.

### *Configuring Port Speed on PTX10K-1201 line card at PIC Level*

To configure all ports to operate at the same speed, configure rate selectability at the PIC level. When you configure rate selectability at the PIC level, all the ports of the PIC that support the configured speed operate at that speed. The PTX10K-LC1201 line card supports port speeds of 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, 100 Gbps, 200 Gbps, and 400 Gbps.

To configure rate selectability at the PIC level:

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

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

For example:

```
[edit]
user@host# edit chassis fpc 5 pic 0
```

2. Configure the **pic-mode *pic-speed*** statement to set the operating speed for the PIC. All ports of the PIC that support the configured speed operate at the configured speed. Values for the ***pic-speed*** option are 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, 100 Gbps, 200 Gbps, and 400 Gbps.

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

For example:

```
[edit chassis fpc 5 pic 0]
user@host# set pic-mode 200G
```

3. Verify the configuration.

```
[edit chassis fpc 5 pic 0]
user@host# show
pic-mode 200G;
```

4. Commit your configuration changes.

In this example, you have configured all ports on the line card with port speed of 200 Gbps.

#### ***Configuring Port Speed on PTX10K-LC1201 line card at Port Level***

To configure different port speeds for each port, you configure rate selectability at the port level. Configuring rate selectability at the port level provides you the flexibility of operating individual ports of the PIC at different supported speeds.

To configure rate selectability at the port level:

1. In configuration mode, navigate to the **[edit chassis fpc *fpc-slot* pic *pic-number*]** hierarchy level.

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

For example:

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

2. To indicate the speed at which the ports operate, configure the **speed** statement for the desired ports. Values supported are 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, 100 Gbps, 200 Gbps, and 400 Gbps. You can also specify the number of sub ports that you wish to configure.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G |25G } 40G | 50G| 100G| 200G } 400G) number-of-sub-
ports port-number
```

For example:

```
[edit chassis fpc 0 pic 0]
user@host# set port 0 speed 10G number-of-sub-ports 4
```

3. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
port 0 {
 speed 10g;
 number-of-sub-ports 4;
}
```

4. Commit your configuration changes.

In this example, you have configured port 0 to operate at 4x10 Gbps speed.

**NOTE:** When you configure the port speed and number of sub ports, 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.

## RELATED DOCUMENTATION

[number-of-ports | 1040](#)

[pic-mode | 1071](#)

[speed | 1134](#)

[Port Speed for Routing Devices | 0](#)

## RELATED DOCUMENTATION

*Introduction to Port Speed*

*Interface Naming Conventions*

### Release History Table

| Release    | Description                                                                                                                                                                                               |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20.4R1 EVO | 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. |
| 20.4R1     | Starting with Junos OS Release 20.4R1, you can now configure 1-Gbps speed on 10-Gigabit Ethernet ports of the MPC7E-10G.                                                                                  |
| 20.1R2     | 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.                                          |
| 19.4R1     | Starting with Junos OS Release 19.4R1, you can now configure 1-Gbps speed on 10-Gigabit Ethernet ports of the JNP10K-LC2101 MPC.                                                                          |
| 18.1R1     | Starting with Junos OS Release 18.1R1, the non-MACsec MIC on the MX10003 routers support 1-Gigabit Ethernet mode also on 10-Gigabit Ethernet mode ports.                                                  |
| 18.1R1     | Starting with Junos OS Release 18.1R1, the 10-Gbps port can operate in 1-Gbps mode also.                                                                                                                  |
| 15.1       | Starting with Junos OS Release 15.1, some PICs support multiple port speeds.                                                                                                                              |

## RELATED DOCUMENTATION

[speed \(Ethernet\) | 1137](#)

[show interfaces port-profile | 1695](#)

## Configuring the Port Speed on the JNP10K-LC1201 by Using New Port Profile Configuration

### IN THIS SECTION

- [Understanding New Port Profile Configuration | 452](#)
- [Configuring Port Speed on JNP10K-LC1201 by using New Port Profile Configuration | 454](#)

This topic describes the JNP10K-LC1201 line card, the guidelines for configuring the port speed and how to configure the port speed using the new port profile configuration.

### Understanding New Port Profile Configuration

#### IN THIS SECTION

- [Guidelines for Configuring the Port Speed Using new Port Profile Configuration | 453](#)

The JNP10K-LC1201 line card is a fixed-configuration line card with 36 built-in ports. The ports on the JNP10K-LC1201 support multiple port speeds. The default port speed is 400Gbps for all ports. Each JNP10K-LC1201 provides a maximum bandwidth of 14.4Tbps.

On the JNP10K-LC1201, you can choose to configure all 36 ports with the following port speeds:

- 40 Gbps, 100 Gbps, 400 Gbps

[Table 94 on page 453](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

Table 94: Rate Selectability for the JNP10K-LC1201

| PIC   | Port Number | Port Speed Supported                                                                                                                                                                                                              | Optics Supported                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 | 0-35        | 4x10-Gigabit Ethernet<br>1x40-Gigabit Ethernet<br>4x25-Gigabit Ethernet<br>8x25-Gigabit Ethernet<br>2x50-Gigabit Ethernet<br>1x100-Gigabit Ethernet<br>2x100-Gigabit Ethernet<br>4x100-Gigabit Ethernet<br>1x400-Gigabit Ethernet | <ul style="list-style-type: none"> <li>• 1x40GE and 4x10GE support using QSFP+</li> <li>• 4x25GE support using QSFP28 25G optics (using breakout cables).</li> <li>• 8x25GE support using QSFP28 25G optics (using breakout cables).</li> <li>• 2x50GE support using QSFP28 50G optics (using breakout cables).</li> <li>• 1x100GE support using QSFP28 100G optics.</li> <li>• 2x100GE support using QSFP28 DD 200G</li> <li>• 4x100GE and 1x400GE support using QSFP56 DD 400G</li> </ul> <p><b>NOTE:</b> By default, all the active ports operate in 400-Gigabit Ethernet mode.</p> |

Starting in Junos OS Evolved Release 20.1R2, we now support a new port profile configuration to configure port speeds on the JNP10K-LC1201 line card. You can now configure the port speed on the JNP10K-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 JNP10K-LC1201 line card.

### Guidelines for Configuring the Port Speed Using new Port Profile Configuration

This topic describes the guidelines to consider when configuring rate selectability using the new port profile configuration:

- 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.
- 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.
- You can configure the active physical ports in a PIC. By default, all the physical ports in an interface are active. To control the number of interfaces created on a physical port in a PIC, use the **unused** command. If you configure a port as unused, no interfaces (channelized or non-channelized) are created for that port.
- 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.
- If you have not configured the port profile for a specific port and the port is active, interfaces are created for the port with default speed based on the platform or FPC. You can control the number of interfaces created by using the **unused** command.
- Interfaces are created 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.

## Configuring Port Speed on JNP10K-LC1201 by using New Port Profile Configuration

In earlier releases, you configured the port speed by using the interface commands in the port profile configuration which was part of the **[edit chassis]** hierarchy. You can now configure the port speed on the JNP10K-LC1201 line card by using the port profile configuration commands in the **[edit interfaces]** hierarchy.

To configure the port speed for non-channelized interfaces on JNP10K-LC1201 :

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

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

For example:

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

2. To indicate the speed at which the ports operate, configure the **speed** statement for the desired interfaces. Values supported are 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, 100 Gbps, and 400 Gbps.

```
[edit interfaces interface-name]
user@host# set speed (10g | 25g | 40g | 50g| 100g | 400g)
```

For example:

```
[edit interfaces et-1/0/3]
user@host# set speed 100g
```

3. (Optional) To control the number of interfaces created on a physical port, use the **unused** statement. If you configure a port as unused, no interfaces are created for that port irrespective of the port profile configuration for that port.

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

For example:

```
[edit]
user@host# set interfaces et-2/0/3 unused
```

No interfaces (channelized or non-channelized) are created on port 3 of the JNP10K-LC1201 line card installed in the FPC slot 2.

To configure the port speed for channelized interfaces on JNP10K-LC1201:

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

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

For example:

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

2. To indicate the speed at which the ports operate, configure the **speed** statement for the desired interfaces. Values supported are 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, 100 Gbps, and 400 Gbps.

```
[edit interfaces interface-name]
user@host# set speed (10g | 25g | 40g | 50g | 100g | 400g)
```

For example:

```
[edit interfaces et-1/0/3]
user@host# set speed 100g
```

3. To specify the number of interfaces per port that you want to configure, use the **number-of-sub-ports** statement.

```
[edit interfaces interface-name]
user@host# set number-of-sub-ports number-of-sub-ports
```

For example:

```
[edit interfaces et-1/0/3]
user@host# set number-of-sub-ports 4
```

In Step 2 and Step 3, you have configured 4x100GE channelized interfaces.

4. (Optional) To control the number of interfaces created on a physical port, use the **unused** statement. If you configure a port as unused, no interfaces are created for that port irrespective of the port profile configuration for that port.

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

For example:

```
[edit]
user@host# set interfaces et-2/0/4 unused
```

No interfaces (channelized or non-channelized) are created on port 4 of the JNP10K-LC1201 line card installed in the FPC slot 2.

## SEE ALSO

[Configuring the Port Speed on the JNP10K-LC1201 by Using New Port Profile Configuration | 452](#)  
[number-of-sub-ports | 1042](#)  
[Unused | 1220](#)

## Release History Table

| Release | Description                                                                                                                                           |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20.1R2  | Starting in Junos OS Evolved Release 20.1R2, we now support a new port profile configuration to configure port speeds on the JNP10K-LC1201 line card. |

# Configuring Optical Transport network

## IN THIS CHAPTER

- Understanding Optical Transport Network (OTN) | 458
- Supported OTN and Optics Options | 491
- Forward Error Correction (FEC) and Bit Error Rate (BER) | 543
- Dense Wavelength Division Multiplexing (DWDM) Interface Wavelength | 552
- Configuring OTN | 556
- ODU Path Delay Measurement for Performance Monitoring | 587

## Understanding Optical Transport Network (OTN)

### IN THIS SECTION

- 10-Gigabit Ethernet OTN Options Configuration Overview | 459
- 100-Gigabit Ethernet OTN Options Configuration Overview | 459
- Understanding the P1-PTX-24-10G-W-SFPP PIC | 461
- Understanding the features of ACX6360 | 466
- Understanding the P2-100GE-OTN PIC | 469
- Understanding the MIC3-100G-DWDM MIC | 474
- Understanding the PTX-5-100G-WDM PIC | 478
- Understanding the PTX10K-LC1104 Line Card | 482
- Interface Mapping and Modulation format for ACX6360 | 485
- Interface Mapping and Modulation format for PTX10K-LC1104 Line Card | 488

Use this topic for overview information about Optical Transport Network support provided by specific line cards and devices.

## 10-Gigabit Ethernet OTN Options Configuration Overview

MX240, MX480, MX960, MX2010, MX2020, T320, T640, T1600, PTX3000, and PTX5000 routers support Optical Transport Network (OTN) interfaces, including the 10-Gigabit Ethernet DWDM OTN PIC, and provide ITU-T G.709 support. Use the `set otn-options` statement at the `[edit interfaces if-fpc/pic/port]` hierarchy level to configure the OTN options.

MX2020, MX2010, MX960, MX480, and MX240 routers support OTN interfaces on MPC5E and MPC6E. MPC5E-40G10G and MPC5EQ-10G40G support OTN on 10-Gigabit Ethernet interfaces but not on 40-Gigabit Ethernet interfaces. The OTN MIC MIC6-10G-OTN on MPC6E supports OTN on 10-Gigabit Ethernet interfaces on MX2020 and MX2010 routers.

**NOTE:** The MIC6-10G-OTN line cards support dual rate (10GE or OTU4). To configure OTN on the MIC6-10G-OTN line card, you must configure the rate on the port to OTU4 by using the `rate rate` statement. If you do not configure the rate to OTU4, OTN mode is not supported on the line card and the interface does not come up.

OTN support on the specified MX Series routers includes:

- International Telecommunications Union (ITU)-standard OTN performance monitoring and alarm management
- Transparent transport of 24 10-Gigabit Ethernet signals with optical channel data unit 2 (ODU2) and ODU2e framing on a per-port basis
- Pre-forward error correction (pre-FEC)-based bit error rate (BER). Fast reroute (FRR) uses the pre-FEC BER as an indication of the condition of an OTN link.

To configure the OTN options on the specified MX routers, use the `set otn-options` statement at the `[edit interfaces interfaceType-fpc/pic/port]` hierarchy level.

### SEE ALSO

| [otn-options](#)

## 100-Gigabit Ethernet OTN Options Configuration Overview

PTX Series routers support optical transport network (OTN) interfaces, including the 100-Gigabit DWDM OTN PIC, which supports:

- Transparent transport of two 100-Gigabit Ethernet signals with Optical Channel Transport Unit 4 (OTU4) framing.
- International Telecommunications Union (ITU)-standard OTN performance monitoring (PM) and alarm management.
- Dual polarization quadrature phase shift keying (DP-QPSK) modulation and soft-decision forward error correction (SD-FEC) for long haul and metro applications.
- Pre-forward error correction (pre-FEC)-based bit error rate (BER) monitoring. Pre-FEC BER monitoring uses the pre-FEC BER as an indication of the condition of an OTN link. See *Understanding Pre-FEC BER Monitoring and BER Thresholds* for more information.

For more information about the 100-Gigabit DWDM OTN PIC, see [100-Gigabit DWDM OTN PIC](#) in the [PTX Series Interface Module Reference](#).

PTX Series routers also support the 100-Gigabit Ethernet OTN PIC (P2-100GE-OTN), which provides four 100-Gigabit Ethernet interfaces, independently configurable in LAN PHY framing mode or in optical channel transport unit 4 (OTU4) mode. See *Understanding the P2-100GE-OTN PIC* for more information.

See *Supported OTN Options on PTX Series Routers* for a comparison of the features supported on PTX Series OTN PICs.

MX2020, MX2010, MX960, MX480, and MX240 routers support OTN interfaces on MPC5E and MPC6E. MPC5E-100G10G and MPC5EQ-100G10G support 100-Gigabit Ethernet OTN interfaces and 10-Gigabit Ethernet OTN interfaces on MX240, MX480, and MX960 routers. The OTN MIC MIC6-100G-CFP2 on MPC6E supports OTN on 100-Gigabit Ethernet interfaces on MX2020 and MX2010 routers. OTN support on the specified MX Series routers includes:

- International Telecommunications Union (ITU)-standard OTN performance monitoring (PM) and alarm management
- Transparent transport of two 100-Gigabit Ethernet signals with optical channel transport unit 4 (OTU4) framing.
- Generic forward error correction (Generic FEC)

To configure the OTN options for PTX Series routers and specific MX Series routers, use the `set otn-options` statement at the `[edit interfaces interfaceType-fpcl picl port]` hierarchy level.

Use the `set optics-options` statement at the `[edit interfaces interfaceType-fpcl picl port]` hierarchy level to configure the optics options.

Use the `show interfaces extensive`, `show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)`, and `show interfaces transport pm` commands to view optics and OTN PM information. To display the current time interval

and clear the channel service unit (CSU) alarm and defect counters, use the *clear interfaces interval* command.

## SEE ALSO

*Configuring OTN Interfaces on P1-PTX-2-100G-WDM*

*show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)*

*optics-options*

*otn-options*

## Understanding the P1-PTX-24-10G-W-SFPP PIC

### IN THIS SECTION

- [Interface Features | 461](#)
- [Layer 2 and Layer 3 Features | 464](#)
- [OTN Alarms and Defects | 465](#)
- [TCA Alarms | 466](#)

Starting from Junos OS Release 14.2, a 24-port 10-Gigabit Ethernet OTN PIC—P1-PTX-24-10G-W-SFPP—is supported on the FPC-PTX-P1-A and FPC2-PTX-P1A FPCs in PTX5000 routers, and the FPC-SFF-PTX-P1-A and FPC-SFF-PTX-T FPCs in PTX3000 routers. The P1-PTX-24-10G-W-SFPP PIC provides twenty-four 10-Gigabit Ethernet interfaces, that are independently configurable in LAN PHY or WAN PHY framing mode or in optical channel transport unit in OTU2e, OTU1e, or OTU2 mode.

The following sections explain this PIC in detail:

### Interface Features

The following interface features are supported on the P1-PTX-24-10G-W-SFPP PIC:

- Twenty-four 10-Gigabit Ethernet interfaces, which are independently configurable in LAN PHY or WAN PHY mode or in OTU2e, OTU1e, or OTU2 signal mode. Each interface is terminated by means of a CFP2 transceiver.
- The interfaces are named with prefix *et*.
- Gigabit Ethernet local loopback.

- Link-level pause frames—You can halt the Ethernet interface from transmitting packets for a configured period of time.
- Interface hold timer and interface damping—You can set the **hold-time** statement (in milliseconds) to damp interface transitions.
- External clock.
- Nonstandard tag protocol identifier (TPID):
  - For each 10-Gigabit Ethernet port, you can configure up to eight TPIDs by using the **tag-protocol-id** statement at the [edit interfaces interface-name gigheter-options ethernet-switch-profile] hierarchy level.
  - The **tag-protocol-id** statement can be configured only on the first port (port 0) of the PIC. If any other (nonzero) port has the **tag-protocol-id** configuration, the Routing Engine registers an error in the system log and the configuration is ignored.
  - The **tag-protocol-id** statement configured on port 0 of the PIC also applies to the rest of the ports on that PIC.
- Generic forward error correction (GFEC), ultra forward error correction (UFEC), enhanced forward error correction (EFEC), and no-FEC modes of operation are supported.
- Diagnostics tools:
  - Line loopback
  - Local loopback
- Fast reroute (FRR)—Based on configurable pre-FEC, bit error rate (BER) is supported and is configured using the **ber-threshold-signal-degrade** statement at the [edit interfaces interface-name otn-options signal-degrade] hierarchy level.
- *jnx-ifotn.mib* and *otn-mib* as defined in RFC 3591. Note that according to Junos OS security standard, configurable parameters are not supported through SNMP. Only the *get* operation is available through SNMP.
- FEC statistics—corrected errors and corrected error ratio.
- OTN payload pseudorandom binary sequence (PRBS) generation and checking by enabling or disabling PRBS with the **prbs** or **no-prbs** statement at the [edit interfaces interface-name otn-options] hierarchy level.
- At the physical interface level, **flexible-ethernet-service**, **ethernet-ccc**, and **ethernet-tcc** encapsulations are supported. For **flexible-ethernet-service** encapsulation, the logical level supports **enet2**, **vlan-ccc**, and **vlan-tcc** encapsulations.

- At the logical interface level **dix**, **vlan-ccc**, and **vlan-tcc** encapsulations are supported.
- SNMP management of the PIC based on RFC 3591, Definitions of Managed Objects for the Optical Interface Type:
  - Set functionality
  - Juniper Networks Black-Link MIB
  - IFOTN MIB
  - Optics MIB
  - FRU MIB
- 15-minute and 1-day performance monitoring and historic statistics.
  - Near-end and far-end performance monitoring
  - Threshold-crossing alerts
  - BER performance monitoring
  - FEC performance monitoring
  - Optical performance monitoring

The following features are not supported on the P1-PTX-24-10G-W-SFPP PIC:

- Source MAC learning for accounting
- MAC policing
- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**
- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation
- Per queue flow control
- Generic framing procedure-framed (GFP-F) mapping modes over OTN
- General communication channel (GCC)
- OTN interface-level Automatic Protection Switching (APS)
- Insertion, monitoring, and display of OTN header overhead byte
- Optical harness support

- Transport interface and state model (GR-1093)
- Trace tone support

### Layer 2 and Layer 3 Features

The following Layer 2 and Layer 3 features are supported on the P1-PTX-24-10G-W-SFPP PIC:

- MAC detect link up and link down based on local fault signal or remote fault signal.
- MAC statistics.
- Flow control.
- MAC oversized packet counters based on default MTU value or user-configured MTU value.
- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Per logical interface source address MAC accounting.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.
- Aggregated Ethernet supports 64 child links that can be configured using the **set chassis aggregated-devices maximum-links** configuration command.
- Maximum of 1024 logical interfaces on an aggregated Ethernet physical interface.
- Support for V LAN tagging, flexible VLAN tagging, and stacked VLAN tagging.
- LACP.
- Link protection.
- 802.3 ah OAM.
- 802.1 ag OAM.
- MPLS FRR.
- SNMP.
- Supports per-VLAN queuing (using Packet Forwarding Engine).

## OTN Alarms and Defects

The following OTN alarms and defects are supported on the P1-PTX-24-10G-W-SFPP PIC:

- LOS—Loss Of Signal
- LOF—Loss Of Frame
- LOM—Loss Of Multiframe
- SSF—Server Signal Failure
- TSF—Trail Signal Fail
- OTU-FEC-DEG—Forward Error Correction Degraded
- OTU-FEC-EXE—Excessive Errors, FEC\_FAIL from the transponder
- OTU-AIS—Alarm Indication Signal or all ones signal
- OTU-BDI—Backward Defect Identification
- OTU-IAE—Incoming Alignment Error
- OTU-BIAE—Backward Incoming Alignment Error
- OTU-TTIM—Destination Access Point Identifier [DAPI], Source Access Point Identifier [SAPI], or both mismatch from expected to received
- OTU-SD—Signal Degrade
- OTU-SF—Signal Fail
- CSF—Client Signal Failure
- ODU-LCK—(ODU lock triggers for PM [path monitoring])
- ODU-AIS—(alarm indication signal or all ones signal)
- ODU-OCI—(open connection error)
- ODU-BDI—(backward defect indication)
- ODU-IAE—(incoming alignment error)
- ODU-DAPI-TTIM—DAPI or DAPI/SAPI mismatch from expected to receive
- ODU-SAPI-TTIM—SAPI or DAPI/SAPI mismatch from expected to receive
- ODU-BEI—Backward Error Indication

- ODU-SSF—Server Signal Fail
- ODU-TSF—Trail Signal Fail
- ODU-SD—Signal Degrade
- ODU-SF—Signal Fail
- OPU-PTM—Payload Type Mismatch

### TCA Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minute interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UAS)

### SEE ALSO

| [Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC | 557](#)

## Understanding the features of ACX6360

### IN THIS SECTION

- [Interface Features | 467](#)
- [OTN Alarms and Defects | 467](#)
- [TCA Alarms | 469](#)

Starting in Junos OS Release 18.2R1, the ACX6360 routers with CFP2-DCO pluggable coherent optics, provide high density long haul OTN transport solution.

The following sections explain the features in detail:

## Interface Features

The following interface features are supported on the ACX6360:

- Compliant with ITU G.709.
- Supports 8 CFP2 DCO optical modules.
- Supports minimum channel spacing of 6.25GHz.
- Ethernet pause frames—You can halt the Ethernet interface from transmitting packets for a configured period of time.
- Soft-decision forward error correction mode (SDFEC)-QPSK-100G, 8QAM-200G and 16QAM-200G modes of operation are supported.
- Diagnostics tools:
  - Line loopback
  - Local loopback
- Fast reroute (FRR)—Based on configurable pre-FEC or configurable Q threshold for signal degrade.
- SNMP management based on RFC 3591, Definitions of Managed Objects for the Optical Interface Type:
  - Black Link MIB—jnx-bl.mib
  - IFOTN MIB—jnx-ifotn.mib
  - Optics MIB—jnx-optics.mib
  - FRU MIB—jnx-fru.mib
- Threshold-crossing alerts
- BER performance monitoring
- FEC performance monitoring
- Optical performance monitoring

## OTN Alarms and Defects

The following OTN alarms and defects are supported on the ACX6360 routers:

- SSF—Server Signal Failure

- TSF—Trail Signal Fail
- OTU-AIS—Alarm Indication Signal or all ones signal
- OTU-BDI—Backward Defect Identification
- OTU-IAE—Incoming Alignment Error
- OTU-BIAE—Backward Incoming Alignment Error
- OTU-TTIM—Destination Access Point Identifier [DAPI], Source Access Point Identifier [SAPI], or both mismatch from expected to received
- OTU-SD—Signal Degrade
- OTU-SSF—Server Signal Fail
- OTU-TSF—Trail Signal Fail
- PRE\_FEC\_SD
- FE\_PRE\_FEC\_SD
- ODU-LCK—(ODU lock triggers for PM [path monitoring])
- ODU-AIS—(alarm indication signal or all ones signal)
- ODU-OCI—(open connection error)
- ODU-BDI—(backward defect indication)
- ODU-IAE—(incoming alignment error)
- ODU-TTIM—DAPI or SAPI mismatch from expected to receive
- ODU-BEI—Backward Error Indication
- ODU-LTC—Loss of tandem connection
- ODU-SSF—Server Signal Fail
- ODU-TSF—Trail Signal Fail
- ODU-CSF—Client Signal Fail
- ODU-SD—Signal Degrade
- ODU-SF—Signal Fail
- OPU-PTM—Payload Type Mismatch

## TCA Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minute interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UAS)

## SEE ALSO

| [Configuring OTN | 556](#)

## Understanding the P2-100GE-OTN PIC

### IN THIS SECTION

- [Interface Features | 470](#)
- [Layer 2 and Layer 3 Features | 472](#)
- [OTN Alarms and Defects | 473](#)
- [TCA Alarms | 474](#)

Starting with Junos OS Release 14.1R2 and 14.2, a 100-Gigabit Ethernet OTN PIC—P2-100GE-OTN—is supported on the FPC2-PTX-P1A FPC in PTX5000 routers. The P2-100GE-OTN PIC provides 4-port 100-Gigabit Ethernet interfaces, which are independently configurable in LAN PHY framing mode or in optical channel transport unit 4 (OTU4) mode. Each interface is terminated by means of a CFP2 transceiver. The FPC2-PTX-P1A FPC supports two P2-100GE-OTN PICs, in which each 100-Gigabit Ethernet port is mapped to a Packet Forwarding Engine in the FPC.

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.

For information about mixed rates, see [Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles](#).

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.

The following sections explain this PIC in detail:

### Interface Features

The following interface features are supported on a P2-100GE-OTN PIC:

- 4-port 100-Gigabit Ethernet interfaces, which are independently configurable in LAN PHY framing mode or in OTU4 signal mode. Each interface is terminated by means of a CFP2 transceiver.
- Each port maps to a single Packet Forwarding Engine in the FPC2-PTX-P1A FPC.
- The interfaces are named with prefix *et*.
- Gigabit Ethernet local loopback.
- Link-level pause frames—You can halt the Ethernet interface from transmitting packets for a configured period of time.
- Interface hold timer and interface damping—You can set the **hold-time** statement (in milliseconds) to damp interface transitions.
- External clock
- Nonstandard tag protocol identifier (TPID):
  - For each 100-Gigabit Ethernet port, you can configure up to eight TPIDs by using the **tag-protocol-id** statement at the **[edit interfaces interface-name gigheter-options ethernet-switch-profile]** hierarchy level.
  - The **tag-protocol-id** statement can be configured only on the first port (port 0) of the PIC. If any other (nonzero) port has the **tag-protocol-id** configuration, the Routing Engine registers an error in the system log and the configuration is ignored.
  - The **tag-protocol-id** statement configured on port 0 of the PIC also applies to the rest of the ports on that PIC.
- The interface *Link Down* event always generates an interrupt; however, the interface *Link Up* event does not generate an interrupt. Therefore, the interface link-up event is detected during the 1-second PIC periodic polling process.
- Generic forward error correction (GFEC) (G.709) and no-FEC modes of operation.
- Diagnostics tools:

- Line loopback
- Local loopback
- Fast reroute (FRR)—Based on configurable pre-FEC, bit error rate (BER) is supported and is configured using the **ber-threshold-signal-degrade** statement at the [edit interfaces *interface-name* otn-options signal-degrade] hierarchy level.
- *jnx-ifotn.mib* and *otn-mib* as defined in RFC 3591. Note that according to Junos OS security standard, configurable parameters are not supported through SNMP. Only the *get* operation is available through SNMP.
- FEC statistics—corrected errors and corrected error ratio.
- OTN payload pseudorandom binary sequence (PRBS) generation and checking by enabling or disabling PRBS with the **prbs** or **no-prbs** statement at the [edit interfaces *interface-name* otn-options] hierarchy level.
- Optical channel data unit (ODU)-level delay measurement.
- At the physical interface level, **flexible-ethernet-service**, **ethernet-ccc**, and **ethernet-tcc** encapsulations are supported. For the **flexible-ethernet-service** encapsulation, the logical level supports **enet2**, **vlan-ccc**, and **vlan-tcc** encapsulations.
- At the logical interface level, **dix**, **vlan-ccc**, and **vlan-tcc** encapsulations are supported.
- Interoperability between 100-Gigabit Ethernet interfaces with CFP transceiver and 100-Gigabit Ethernet interfaces with CFP2 transceiver in LAN PHY framing mode and in OTU4 mode.

The following features are not supported on the P2-100GE-OTN PIC:

- Source MAC learning for accounting
- MAC policing
- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**
- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation
- Per-queue flow control
- Generic framing procedure-framed (GFP-F) mapping modes over OTN
- General communication channel (GCC)
- OTN interface-level Automatic Protection Switching (APS)

- Insertion, monitoring, and display of OTN header overhead byte
- Black link MIB for integration with transponders
- Optical harness support
- Transport interface and state model (GR-1093)
- Trace tone support
- 15-minute and 1-day performance monitoring counters and historic counters

### Layer 2 and Layer 3 Features

The following Layer 2 and Layer 3 features are supported on the P2-100GE-OTN PIC:

- MAC detect link up and link down based on local fault signal or remote fault signal.
- MAC statistics.
- Flow control.
- MAC oversized packet counters based on default MTU value or user-configured MTU value.
- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Per-logical interface source address MAC accounting.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.
- Aggregated Ethernet supports 64 child links that can be configured using the **set chassis aggregated-devices maximum-links** configuration command.
- Maximum of 1024 logical interfaces on an aggregated Ethernet physical interface.
- Support for VLAN tagging, flexible VLAN tagging, and stacked VLAN tagging.
- LACP.
- Link protection.
- 802.3 ah OAM.
- 802.1 ag OAM.

- MPLS FRR.
- SNMP.
- Supports per-VLAN queuing (using Packet Forwarding Engine).

### **OTN Alarms and Defects**

The following OTN alarms and defects are supported on the P2-100GE-OTN PIC:

- LOS—Loss Of Signal
- LOF—Loss Of Frame
- LOM—Loss Of Multiframe
- OTU—Degrade
- OTU—AIS
- OTU—IAE
- OTU—BDI
- OTU—TTIM
- OTU—Signal Degrade
- OTU—Signal Fail
- ODU—Signal Fail
- OTU-FEC—Degrade
- OTU-FEC—Excessive errors
- ODU—Signal Degrade
- ODU—AIS
- ODU—BDI
- ODU—OCI
- ODU—LCK
- ODU—TTIM
- OPU—PTM

## TCA Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minute interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UAS)

## SEE ALSO

| [Configuring OTN Interfaces on P2-100GE-OTN PIC](#) | 563

## Understanding the MIC3-100G-DWDM MIC

### IN THIS SECTION

- [Interface Features](#) | 475
- [Layer 2 and Layer 3 Features](#) | 476
- [OTN Alarms and Defects](#) | 476

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. The MIC3-100G-DWDM MIC provides a single 100-Gigabit Ethernet interface port that supports DP-QPSK with coherent reception and OTU4 and OTU4 (v) framing modes.

The interfaces on MIC3-100G-DWDM MIC are named with prefix *et*. For more information, see [Interface Naming Overview](#).

The following sections explain the features of this MIC in detail:

## Interface Features

The following interface features are supported on the MIC3-100G-DWDM MIC:

- Single port 100-Gigabit Ethernet interface with OTU4 (v) framing. DP-QPSK modulation with coherent reception using a CFP2-ACO DWDM optical transceiver.
- Gigabit Ethernet local loopback.
- Diagnostics tools:
  - Line loopback
  - Local loopback
  - Optical Channel Data Unit (ODU) Open Connection Error
  - Optical Channel Data Unit (ODU) Lock Maintenance Signal
- Types of forward error corrections (FEC):
  - GFEC (generic forward error correction)
  - HGFECC (high gain forward error correction)
  - SDFEC (soft-decision forward error correction)
- The following MIB modules continue to be supported (and no new MIB is introduced):
  - MIB module to describe Black Link extension to RFC 3591 (jnxoptIfExtMibModule)
  - MIB module to manage the OTN interface (jnxIfOtnMib)
  - MIB module to manage the Optics interface (jnxIfOpticsMib)
  - MIB module to manage OTN FRUs (jnxFruMib)
- Interoperability with the 100-Gigabit DWDM OTN PIC (P1-PTX-2-100G-WDM) is not supported.
- Support for interoperability with other vendors' 100 Gigabit Ethernet interfaces.
- Source MAC learning for accounting
- MAC policing
- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**
- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation

- Per-queue flow control
- 15-minute and 1-day performance monitoring and historic statistics.
  - Near-end and far-end performance monitoring
  - Threshold-crossing alarms
  - BER performance monitoring
  - FEC performance monitoring
  - Optical performance monitoring
- Insertion, monitoring, and display of OTN header overhead
- Transport interface and state model (GR-1093)

### **Layer 2 and Layer 3 Features**

The following Layer 2 and Layer 3 features are supported on the MIC3-100G-DWDM MIC:

- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.
- Flexible VLAN tagging.
- 802.3 ah OAM.
- 802.1 ag OAM.

### **OTN Alarms and Defects**

The following OTN alarms and defects are supported on the MIC3-100G-DWDM MIC:

Optical Channel(OC) Alarms and Defects

- OC-LOS—Loss Of Signal
- OC-LOF—Loss Of Frame
- OC-LOM—Loss Of Multiframe

- OC-Wavelength-Lock—Wavelength Lock

#### Optical Channel Data Unit (ODU) Defects

- ODU-AIS—ODU Alarm Indication Signal
- ODU-BDI—ODU Backward Defect Indication
- ODU-BIAE—ODU Backward Incoming Alignment Error
- ODU-IAE—ODU Incoming Alignment Error
- ODU-LCK—ODU Locked
- ODU-LTC—ODU Loss of Tandem Connection
- ODU-OCI—ODU Open Connection Error
- ODU-SSF—ODU Server Signal Failure
- ODU-TSF—ODU Trail Signal Failure
- ODU-TTIM—ODU Trail Trace Identifier Mismatch

#### Optical Channel Transport Unit (OTU) Defects

- OTU-AIS—OTU Alarm Indication Signal
- OTU-BDI—OTU Backward Defect Indication
- OTU-BIAE—OTU Backward Incoming Alignment Error
- OTU-FEC-DEG—OTU Forward Error Correction Degrade
- OTU-FEC-EXCESS-FEC—OTU Forward Error Correction Excessive FEC Errors
- OTU-IAE—OTU Incoming Alignment Error
- OTU-SSF—OTU Server Signal Failure
- OTU-TSF—OTU Trail Signal Failure
- OTU-TTIM—OTU Trail Trace Identifier Mismatch

#### Threshold-Crossing Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold — near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minutes interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UES)

## SEE ALSO

*Before You Begin Installing or Upgrading the Firmware*

[Configuring OTN Interfaces on MIC3-100G-DWDM MIC | 569](#)

[Configuring Packet Optical Networks with PTX Series Devices](#)

## Understanding the PTX-5-100G-WDM PIC

### IN THIS SECTION

- [Interface Features | 479](#)
- [Layer 2 and Layer 3 Features | 480](#)
- [OTN Alarms and Defects | 480](#)

Starting with Junos OS Release 15.1F6, the 5-port 100-Gigabit DWDM OTN PIC—PTX-5-100G-WDM—is supported on the PTX3000 and the PTX5000 routers. The PTX-5-100G-WDM PIC provides five 100-Gigabit Ethernet interface ports that support dual-polarization quadrature phase shift keying (DP-QPSK) modulation with coherent reception and OTU4 and OTU4 (v) framing modes.

The interfaces on the PTX-5-100G-WDM PIC are named with the prefix *et*. For more information, see [Interface Naming Overview](#).

**NOTE:** The 5-port 100-Gigabit DWDM OTN PIC is not directly interoperable with the 2-port 100-Gigabit DWDM OTN PIC (P1-PTX-2-100G-WDM), but they can both operate over the same DWDM line system.

The following sections explain the features of this PIC in detail:

## Interface Features

The following interface features are supported on the PTX-5-100G-WDM PIC:

- Five-port 100-Gigabit Ethernet interface with OTU4 (v) framing and DP-QPSK modulation with coherent reception using a CFP2-ACO DWDM optical transceiver.
- Gigabit Ethernet local loopback.
- Diagnostics tools:
  - Line loopback
  - Local loopback
  - Optical Channel Data Unit (ODU) Open Connection Error
  - Optical Channel Data Unit (ODU) Lock Maintenance Signal
- Types of forward error corrections (FEC):
  - GFEC (generic forward error correction)

**NOTE:** GFEC mode is not supported on Junos OS Release 15.1F6. Junos OS Release 15.1F6-S1 supports GFEC mode. Contact customer support for the Junos OS Release 15.1F6-S1.

- SDFEC (soft-decision forward error correction)
- The following MIB features continue to be supported (and no new MIB is introduced):
  - MIB module to describe Black Link extension to RFC 3591 (jnxoptIfExtMibModule). The Black Link extension enables an optical transceiver of a vendor to introduce an optical signal over an optical network from another vendor.
  - MIB module to manage the OTN interface (jnxIfOtnMib)
  - MIB module to manage the Optics interface (jnxIfOpticsMib)
  - MIB module to manage OTN FRUs (jnxFruMib)
- Interoperability with other vendors' 100 Gigabit-Ethernet interfaces.
- Source MAC learning for accounting
- MAC policing
- Physical interface-level encapsulations—**vlan-ccc**, **extended-vlan-ccc**, and **extended-vlan-tcc**

- Logical interface-level encapsulation—**vlan-vpls**
- VLAN rewrite for **ccc** encapsulation
- Per-queue flow control
- 15-minute and 1-day performance monitoring and historic statistics.
  - Near-end and far-end performance monitoring
  - Threshold-crossing alarms
  - BER performance monitoring
  - FEC performance monitoring
  - Optical performance monitoring
- Insertion, monitoring, and display of OTN header overhead
- Transport interface and state model (GR-1093)

### **Layer 2 and Layer 3 Features**

The following Layer 2 and Layer 3 features are supported on the PTX-5-100G-WDM PIC:

- Per-port destination address MAC filter.
- Per-port source address MAC filter.
- Per-physical interface source address MAC filter.
- Maximum of 1000 source MAC filter per physical interface.
- Maximum of 32,000 filter terms to share across all filter features.
- Flexible VLAN tagging.
- 802.3 ah OAM.
- 802.1 ag OAM.

### **OTN Alarms and Defects**

The following OTN alarms and defects are supported on the PTX-5-100G-WDM PIC:

#### Optical Channel Alarms and Defects

- OC-LOS—Loss Of Signal

- OC-LOF—Loss Of Frame
- OC-LOM—Loss Of Multiframe
- OC-Wavelength-Lock—Wavelength Lock

#### Optical Channel Data Unit (ODU) Defects

- ODU-AIS—ODU Alarm Indication Signal
- ODU-BDI—ODU Backward Defect Indication
- ODU-BIAE—ODU Backward Incoming Alignment Error
- ODU-IAE—ODU Incoming Alignment Error
- ODU-LCK—ODU Locked
- ODU-LTC—ODU Loss of Tandem Connection
- ODU-OCI—ODU Open Connection Error
- ODU-SSF—ODU Server Signal Failure
- ODU-TSF—ODU Trail Signal Failure
- ODU-TTIM—ODU Trail Trace Identifier Mismatch

#### Optical Channel Transport Unit (OTU) Defects

- OTU-AIS—OTU Alarm Indication Signal
- OTU-BDI—OTU Backward Defect Indication
- OTU-BIAE—OTU Backward Incoming Alignment Error
- OTU-FEC-DEG—OTU Forward Error Correction Degrade
- OTU-FEC-EXCESS-FEC—OTU Forward Error Correction Excessive FEC Errors
- OTU-IAE—OTU Incoming Alignment Error
- OTU-SSF—OTU Server Signal Failure
- OTU-TSF—OTU Trail Signal Failure
- OTU-TTIM—OTU Trail Trace Identifier Mismatch

#### Threshold Crossing Alarms

Threshold-crossing alarms (TCAs) are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15-minute interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UES)

## SEE ALSO

*Before You Begin Installing or Upgrading the Firmware*

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 575](#)

*Installing Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*

*Upgrading Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*

[Configuring Packet Optical Networks with PTX Series Devices](#)

## Understanding the PTX10K-LC1104 Line Card

### IN THIS SECTION

- [Software Features | 482](#)
- [OTN Alarms and Defects | 484](#)

The PTX10K-LC1104 line card provides up to 1.2 Tbps packet forwarding for cloud providers, service providers, and enterprises that need coherent dense wavelength-division multiplexing (DWDM) with MACsec security features.

The PTX10K-LC1104 line card is supported on Junos OS Release 18.3R1 and later.

The following sections explain the features of the PTX10K-LC1104 line card in detail:

### Software Features

The following interface features are supported on the PTX10K-LC1104:

- Compliant with ITU G.709 and G.798

- Performance monitoring features such as alarms, threshold-crossing alarms, OTU/ODU error seconds, and FEC and bit error rate (BER) statistics.
- SNMP management of the MIC based on RFC 3591, Managed Objects for the Optical Interface Type, including the following:
  - Black Link MIB-jnx-bl.mib
  - IFOTN MIB-jnx-ifotn.mib
  - Optics MIB-jnx-optics.mib
  - FRU MIB-jnx-fru.mib
- User-configurable optics options:
  - Modulation format: 16QAM, 8QAM, QPSK
  - FEC mode (15% SDFEC or 25% SDFEC)
  - Differential and non-differential encoding modes
  - Transmit (TX) laser enable and disable
  - TX output power
  - Wavelength
  - Threshold crossing alarms (TCAs)
- IEEE 802.1ag OAM
- IEEE 802.3ah OAM
- IFINFO/IFMON
- IEEE 802.3ad link aggregation
- Flexible Ethernet services encapsulation
- Flexible VLAN tagging
- Source address MAC accounting per logical interface
- Source address MAC filter per port
- Source address MAC filter per logical interface
- Destination address MAC filter per port
- Up to 8000 logical interfaces shared across all ports on a single PFE

## OTN Alarms and Defects

The following OTN alarms and defects are supported on the PTX10K-LC1104 line card:

### Optical Channel(OC) Alarms and Defects

- OC-LOS—Loss Of Signal
- OC-LOF—Loss Of Frame
- OC-LOM—Loss Of Multiframe
- OC-Wavelength-Lock—Wavelength Lock

### Optical Channel Data Unit (ODU) Defects

- ODU-AIS—ODU Alarm Indication Signal
- ODU-BDI—ODU Backward Defect Indication
- ODU-IAE—ODU Incoming Alignment Error
- ODU-LCK—ODU Locked
- ODU-LTC—ODU Loss of Tandem Connection
- ODU-OCI—ODU Open Connection Error
- ODU-SSF—ODU Server Signal Failure
- ODU-TSF—ODU Trail Signal Failure
- ODU-TTIM—ODU Trail Trace Identifier Mismatch

### Optical Channel Transport Unit (OTU) Defects

- OTU-AIS—OTU Alarm Indication Signal
- OTU-BDI—OTU Backward Defect Indication
- OTU-BIAE—OTU Backward Incoming Alignment Error
- OTU-FEC-DEG—OTU Forward Error Correction Degrade
- OTU-FEC-EXCESS-FEC—OTU Forward Error Correction Excessive FEC Errors
- OTU-IAE—OTU Incoming Alignment Error
- OTU-SSF—OTU Server Signal Failure
- OTU-TSF—OTU Trail Signal Failure

- OTU-TTIM—OTU Trail Trace Identifier Mismatch

#### Threshold-Crossing Alarms

Threshold-crossing alarms (TCA) are alarms that are activated when a certain configurable threshold — near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15 minutes interval for parameters such as OTU and ODU. The following alarms are supported:

- Background block error threshold (BBE)
- Errored seconds threshold (ES)
- Severely errored seconds threshold (SES)
- Unavailable seconds threshold (UES)

#### SEE ALSO

[Configuring OTN Interface Options on PTX10K-LC1104 | 581](#)

### Interface Mapping and Modulation format for ACX6360

ACX6360 routers supports 8 CFP2-DCO optical modules. For each CFP2-DCO optical module, 1 *ot* interface is created. Hence, 8 *ot* interfaces are created for ACX6360 routers. ACX6360 routers support only 100GE *et* interfaces and up to 2 *et* interfaces can be mapped to 1 *ot* interface, depending on the configured CFP2-DCO rate- 100G or 200G.

The optical interface to *et* mapping is displayed in the following table:

| "ot-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
| ot-0/1/0        | QPSK-100G         | et-0/1/0                 |
|                 | 8QAM-200G         | et-0/1/0                 |
|                 |                   | et-0/1/1                 |
| 16QAM-200G      | et-0/1/0          |                          |
|                 | et-0/1/1          |                          |

*(Continued)*

| "ot-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
| ot-0/1/1        | QPSK-100G         | et-0/1/2                 |
|                 | 8QAM-200G         | et-0/1/2<br>et-0/1/3     |
|                 | 16QAM-200G        | et-0/1/2<br>et-0/1/3     |
| ot-0/1/2        | QPSK-100G         | et-0/1/4                 |
|                 | 8QAM-200G         | et-0/1/4<br>et-0/1/5     |
|                 | 16QAM-200G        | et-0/1/4<br>et-0/1/5     |
| ot-0/1/3        | QPSK-100G         | et-0/1/6                 |
|                 | 8QAM-200G         | et-0/1/6<br>et-0/1/7     |
|                 | 16QAM-200G        | et-0/1/6<br>et-0/1/7     |
| ot-0/1/4        | QPSK-100G         | et-0/1/8                 |

*(Continued)*

| "ot-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
|                 | 8QAM-200G         | et-0/1/8<br>et-0/1/9     |
|                 | 16QAM-200G        | et-0/1/8<br>et-0/1/9     |
| ot-0/1/5        | QPSK-100G         | et-0/1/10                |
|                 | 8QAM-200G         | et-0/1/10<br>et-0/1/11   |
|                 | 16QAM-200G        | et-0/1/10<br>et-0/1/11   |
| ot-0/1/6        | QPSK-100G         | et-0/1/12                |
|                 | 8QAM-200G         | et-0/1/12<br>et-0/1/13   |
|                 | 16QAM-200G        | et-0/1/12<br>et-0/1/13   |
| ot-0/1/7        | QPSK-100G         | et-0/1/14                |
|                 | 8QAM-200G         | et-0/1/14<br>et-0/1/15   |
|                 |                   | et-0/1/15                |

*(Continued)*

| "ot-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
|                 | 16QAM-200G        | et-0/1/14                |
|                 |                   | et-0/1/15                |

**SEE ALSO**

| [Configuring OTN | 556](#)

**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<br>et-x/0/1     |
|                 | 16QAM             | et-x/0/0<br>et-x/0/1     |
| ot-0/0/1        | QPSK              | et-x/0/2                 |
|                 | 8QAM              | et-x/0/1<br>et-x/0/2     |

*(Continued)*

| "ot-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
|                 | 16QAM             | et-x/0/2<br>et-x/0/3     |
| ot-0/0/2        | QPSK              | et-x/0/4                 |
|                 | 8QAM              | et-x/0/4<br>et-x/0/5     |
|                 | 16QAM             | et-x/0/4<br>et-x/0/5     |
| ot-0/0/3        | QPSK              | et-x/0/6                 |
|                 | 8QAM              | et-x/0/5<br>et-x/0/6     |
|                 | 16QAM             | et-x/0/6<br>et-x/0/7     |
| ot-0/0/4        | QPSK              | et-x/0/8                 |
|                 | 8QAM              | et-x/0/8<br>et-x/0/9     |
|                 | 16QAM             | et-x/0/8<br>et-x/0/9     |

*(Continued)*

| "ot-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
| ot-0/0/5        | QPSK              | et-x/0/10                |
|                 | 8QAM              | et-x/0/9<br>et-x/0/10    |
|                 | 16QAM             | et-x/0/10<br>et-x/0/11   |

**SEE ALSO**

| [Configuring OTN | 556](#)

**Release History Table**

| Release | Description                                                                                                                                                                                                                   |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 | 556](#)

| [Forward Error Correction \(FEC\) and Bit Error Rate \(BER\) | 543](#)

## Supported OTN and Optics Options

### IN THIS SECTION

- Supported OTN Options on PTX Series Routers | [491](#)
- Supported OTN Options on MX Series Routers | [503](#)
- Supported OTN Options on ACX6360 Routers | [515](#)
- Supported OTN Options on ACX5448-D Routers | [519](#)
- Supported OTN Options on PTX10008 and PTX10016 Series Routers | [524](#)
- Supported Optics Options on ACX6360 and ACX5448-D Routers | [533](#)
- Supported Optics Options on PTX10008 and PTX10016 Series Routers | [537](#)

Use this topic for information about the supported optics options and otn options on specific devices.

### Supported OTN Options on PTX Series Routers

[Table 95 on page 491](#) 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. Note that the term *NA* denotes that the statement is not applicable for that particular component:

**Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers**

| Statement           | Options                            | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|---------------------|------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
| bytes (otn-options) | transmit-payload-type <i>value</i> | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement                                                   | Options                                              | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000)              | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-------------------------------------------------------------|------------------------------------------------------|------------------------------------------|----------------------------------------|-------------------------------------------|
| <code>fec</code>                                            | <i>(efec / gfec / gfec-<br/>sdfec / none / ufec)</i> | 13.2/13.3                                | 14.1R2<br><br>14.2<br><br>(gfec, none) | 14.2                                      |
| <b>insert-odu-lck</b>                                       | -                                                    | 13.2/13.3                                | 14.1R2<br><br>14.2                     | 14.2                                      |
| <b>insert-odu-oci</b>                                       | -                                                    | 13.2/13.3                                | 14.1R2<br><br>14.2                     | 14.2                                      |
| <code>is-ma   no-<br/>is-ma</code>                          | -                                                    | 13.2/13.3                                | NA                                     | 14.2                                      |
| <code>laser-<br/>enable   no-<br/>laser-<br/>enable</code>  | -                                                    | 13.2/13.3                                | 14.1R2<br><br>14.2                     | 14.2                                      |
| <code>line-<br/>loopback  <br/>no-line-<br/>loopback</code> | -                                                    | 13.2/13.3                                | 14.1R2<br><br>14.2                     | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement                                     | Options                                             | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------------------------------------------|-----------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
| local-loopback  <br><b>no</b> -local-loopback | -                                                   | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
| odu-delay-management                          | bypass   <b>no</b> -bypass                          | 13.2/13.3                                | NA                        | NA                                        |
|                                               | monitor-end-point   <b>no</b> -monitor-end-point    | 13.2/13.3                                | NA                        | NA                                        |
|                                               | number-of-frames <i>value</i>                       | 13.2/13.3                                | NA                        | NA                                        |
|                                               | <b>no</b> -start-measurement  <br>start-measurement | 13.2/13.3                                | NA                        | NA                                        |
| <b>odu-signal-degrade</b>                     | ber-threshold -<br>clear <i>value</i>               | NA<br>NA                                 | 14.1R2<br><br>14.2        | NA                                        |
|                                               | ber-threshold -<br>signal-degrade<br><i>value</i>   | NA<br>NA                                 | 14.1R2<br><br>14.2        | NA                                        |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement                                                       | Options                                                    | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------------------------------------------------------------|------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|                                                                 | <code>interval <i>value</i></code>                         | NA                                       | 14.1R2<br>14.2            | NA                                        |
| <code>odu-ttim-action-enable   no-odu-ttim-action-enable</code> | -                                                          | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
| <code>otu-ttim-action-enable   no-otu-ttim-action-enable</code> | -                                                          | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
| <code>prbs   no-prbs</code>                                     | -                                                          | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
| <code>preemptive-fast-reroute</code>                            | <code>backward-frr -enable   no-backward-frr-enable</code> | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement | Options                                                                                             | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------|-----------------------------------------------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|           | signal- degrade-<br>monitor- enable  <br><b>no</b> -signal-<br>degrade-<br>monitor-enable           | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | odu-backward-<br>frr-enable   <b>no</b> -<br>odu-backward-<br>frr-enable                            | NA                                       | 14.1R2<br><br>14.2        | NA                                        |
|           | odu-signal-<br>degrade-<br>monitor-enable  <br><b>no</b> -odu-signal-<br>degrade-<br>monitor-enable | NA                                       | 14.1R2<br><br>14.2        | NA                                        |
| rate      | fixed-stuff-<br>bytes   <b>no</b> -fixed-<br>stuff-bytes                                            | 13.2/13.3                                | NA                        | 14.2                                      |
|           | <b>oc192</b>                                                                                        | 13.2/13.3                                | NA                        | 14.2                                      |
|           | otu4                                                                                                | 13.2/13.3                                | 14.1R2<br><br>14.2        | NA                                        |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement      | Options                                                                            | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|----------------|------------------------------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|                | pass-through   no-pass-through                                                     | 13.2/13.3                                | NA                        | NA                                        |
| signal-degrade | ber-threshold-clear <i>value</i>                                                   | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
|                | ber-threshold-signal-degrade <i>value</i>                                          | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
|                | interval <i>value</i>                                                              | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
| tca            | odu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold  <br>threshold-24hrs)    | 13.2/13.3                                | 14.2                      | 14.2                                      |
|                | odu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold  <br>threshold-24hrs) | 13.2/13.3                                | NA                        | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement | Options                                                                          | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------|----------------------------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|           | <b>odu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>     | 13.2/13.3                                | 14.2                      | 14.2                                      |
|           | <b>odu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>  | 13.2/13.3                                | NA                        | 14.2                                      |
|           | <b>odu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 13.2/13.3                                | 14.2                      | 14.2                                      |
|           | <b>odu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 13.2/13.3                                | NA                        | 14.2                                      |
|           | <b>odu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 13.2/13.3                                | 14.2                      | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement | Options                                                                                        | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------|------------------------------------------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|           | <b>otu-tca-bbe</b><br>(enable-tca   no-<br>enable-tca  <br>threshold  <br>threshold-24hrs)     | 13.2/13.3                                | 14.2                      | 14.2                                      |
|           | <b>otu-tca-bbe-fe</b><br>(enable-tca   no-<br>enable-tca  <br>threshold  <br>threshold-24hrs)  | 13.2/13.3                                | NA                        | 14.2                                      |
|           | <b>otu-tca-es</b> (enable-<br>tca   no-enable-tca  <br>threshold  <br>threshold-24hrs)         | 13.2/13.3                                | 14.2                      | 14.2                                      |
|           | <b>otu-tca-es-fe</b><br>(enable-tca   no-<br>enable-tca  <br>threshold  <br>threshold-24hrs)   | 13.2/13.3                                | NA                        | 14.2                                      |
|           | <b>otu-tca-fec-ber</b><br>(enable-tca   no-<br>enable-tca  <br>threshold  <br>threshold-24hrs) | 13.2/13.3                                | 14.2                      | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement                         | Options                                                                          | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------------------------------|----------------------------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|                                   | <b>otu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 13.2/13.3                                | 14.2                      | 14.2                                      |
|                                   | <b>otu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 13.2/13.3                                | NA                        | 14.2                                      |
|                                   | <b>otu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 13.2/13.3                                | 14.2                      | 14.2                                      |
|                                   | <b>otu-tca-uas-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 13.2/13.3                                | NA                        | 14.2                                      |
| <i>transport-monitoring</i>       | -                                                                                | NA                                       | NA                        | 14.2                                      |
| <i>trigger-trigger-identifier</i> | <b>oc-lof (hold-time (down   up)   ignore)</b>                                   | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers *(Continued)*

| Statement | Options                                                            | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------|--------------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|           | <b>oc-lom (hold-time<br/>(down   up)   ignore)</b>                 | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>oc-los (hold-time<br/>(down   up)   ignore)</b>                 | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>oc-tsf (hold-time<br/>(down   up)   ignore)</b>                 | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>oc-wavelength-lock<br/>(hold-time (down  <br/>up)   ignore)</b> | 13.2/13.3                                | NA                        | 14.2                                      |
|           | <b>odu-ais (hold-time<br/>(down   up)   ignore)</b>                | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>odu-bdi (hold-time<br/>(down   up)   ignore)</b>                | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>odu-bei (hold-time<br/>(down   up)   ignore)</b>                | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>odu-iae (hold-time<br/>(down   up)   ignore)</b>                | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement | Options                                                     | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------|-------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|           | <b>odu-lck (hold-time<br/>(down   up)   ignore)</b>         | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>odu-oci (hold-time<br/>(down   up)   ignore)</b>         | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>odu-sd (hold-time<br/>(down   up)   ignore)</b>          | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>odu-ttim</b>                                             | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>opu-ptim (hold-time<br/>(down   up)  <br/>ignore)</b>    | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>otu-ais (hold-time<br/>(down   up)   ignore)</b>         | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>otu-bdi (hold-time<br/>(down   up)   ignore)</b>         | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |
|           | <b>otu-fec-deg (hold-time<br/>(down   up)  <br/>ignore)</b> | 13.2/13.3                                | 14.1R2<br><br>14.2        | 14.2                                      |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement                 | Options                                             | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-<br>PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|---------------------------|-----------------------------------------------------|------------------------------------------|---------------------------|-----------------------------------------------|
|                           | <b>otu-fec-exe (hold-time (down   up)   ignore)</b> | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |
|                           | <b>otu-iae (hold-time (down   up)   ignore)</b>     | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |
|                           | <b>otu-sd (hold-time (down   up)   ignore)</b>      | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |
|                           | <b>otu-ttim (hold-time (down   up)   ignore)</b>    | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |
| <i>tti tti-identifier</i> | <b>odu-dapi <i>identifier</i></b>                   | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |
|                           | <b>odu-expected-receive-dapi <i>identifier</i></b>  | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |
|                           | <b>odu-expected-receive-sapi <i>identifier</i></b>  | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |
|                           | <b>odu-sapi <i>identifier</i></b>                   | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                          |

Table 95: Statements Supported on 100-Gigabit Ethernet PICs on PTX Series Routers (Continued)

| Statement | Options                                                     | P1-PTX-2-100G-WDM<br>(PTX5000 / PTX3000) | P2-100GE-OTN<br>(PTX5000) | P1-PTX-24-10<br>G-W-SFPP<br><br>(PTX5000) |
|-----------|-------------------------------------------------------------|------------------------------------------|---------------------------|-------------------------------------------|
|           | <b>otu-dapi <i>identifier</i></b>                           | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
|           | <b>otu-expected-<br/>receive-dapi<br/><i>identifier</i></b> | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
|           | <b>otu-expected-<br/>receive-sapi<br/><i>identifier</i></b> | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |
|           | <b>otu-sapi <i>identifier</i></b>                           | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                                      |

**SEE ALSO**

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

**Supported OTN Options on MX Series Routers**

Table 96 on page 504 lists the statements that are supported on 100-Gigabit Ethernet MICs on MX Series 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 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers

| Statement                                             | Options                                                               | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-------------------------------------------------------|-----------------------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
| bytes                                                 | <b>transmit-payload-type</b><br><i>value</i>                          | NA                                      | 15.1F5                                                             |
| fec                                                   | <i>(efec / gfec / gfec-sdfec /<br/>hgfec/ none / dsfec/<br/>ufec)</i> | 13.3R3<br><br>(ufec)                    | 15.1F5<br><br>(gfec,hgfec,sdfec)                                   |
| <b>insert- odu-lck</b>                                | -                                                                     | 13.3R3                                  | 15.1F5                                                             |
| <b>insert- odu-oci</b>                                | -                                                                     | 13.3R3                                  | 15.1F5                                                             |
| is-ma   no-is-<br>ma                                  | -                                                                     | 13.3R3                                  | 15.1F5                                                             |
| laser-enable  <br><b>no-laser-</b><br>enable          | -                                                                     | 13.3R3                                  | 15.1F5                                                             |
| line-loopback<br>  <b>no-line-</b><br>loopback        | -                                                                     | 13.3R3                                  | 15.1F5                                                             |
| local-<br>loopback   <b>no-</b><br>local-<br>loopback | -                                                                     | 13.3R3                                  | 15.1F5                                                             |
| odu-delay-<br>management                              | bypass   <b>no-bypass</b>                                             | NA                                      | 15.1F5                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement                                                                      | Options                                                  | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|--------------------------------------------------------------------------------|----------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
|                                                                                | monitor- end-point  <br><b>no</b> -monitor-end-point     | NA                                      | 15.1F5                                                             |
|                                                                                | number- of-frames<br><b>value</b>                        | NA                                      | 15.1F5                                                             |
|                                                                                | <b>no</b> -start-<br>measurement   start-<br>measurement | NA                                      | 15.1F5                                                             |
| <b>signal-degrade</b>                                                          | ber-threshold -clear<br><b>value</b>                     | 13.3R3                                  | 15.1F5                                                             |
|                                                                                | ber-threshold -<br>signal-degrade <b>value</b>           | 13.3R3                                  | 15.1F5                                                             |
|                                                                                | interval <b>value</b>                                    | 13.3R3                                  | 15.1F5                                                             |
| odu-ttim-<br>action-<br>enable   <b>no</b> -<br>odu-ttim-<br>action-<br>enable | -                                                        | 13.3R3                                  | 15.1F5                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement                                                              | Options                                                                              | MIC6-100G-CFP2<br>(MX2010 / MX2020) | MIC3-100G-DWDM<br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|------------------------------------------------------------------------|--------------------------------------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
| otu-ttim-<br>action-<br>enable   no-<br>otu-ttim-<br>action-<br>enable | -                                                                                    | 13.3R3                              | 15.1F5                                                         |
| prbs   no-prbs                                                         | -                                                                                    | 13.3R3                              | 15.1F5                                                         |
| preemptive-<br>fast- reroute                                           | backward-frr -enable<br>  no-backward-frr-<br>enable                                 | 13.3R3                              | 15.1F5                                                         |
|                                                                        | signal- degrade-<br>monitor- enable   no-<br>signal- degrade-<br>monitor-enable      | 13.3R3                              | 15.1F5                                                         |
|                                                                        | odu-backward-frr-<br>enable   no-odu-<br>backward-frr-enable                         | NA                                  | 15.1F5                                                         |
|                                                                        | odu-signal-degrade-<br>monitor-enable   no-<br>odu-signal-degrade-<br>monitor-enable | NA                                  | 15.1F5                                                         |
| rate                                                                   | fixed-stuff-bytes  <br>no-fixed-stuff-bytes                                          | 13.3R3                              | 15.1F5                                                         |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement      | Options                                                                                   | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|----------------|-------------------------------------------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
|                | oc192                                                                                     | 13.3R3                                  | 15.1F5                                                             |
|                | otu4                                                                                      | 13.3R3                                  | 15.1F5                                                             |
|                | pass-through                                                                              | 13.3R3<br><br>(pass-through)            | 15.1F5<br><br>(pass-through)                                       |
| signal-degrade | ber-threshold-clear<br><b>value</b>                                                       | 13.3R3                                  | 15.1F5                                                             |
|                | ber-threshold-signal-degrade<br><b>value</b>                                              | 13.3R3                                  | 15.1F5                                                             |
|                | interval <b>value</b>                                                                     | 13.3R3                                  | 15.1F5                                                             |
|                | <b>q-threshold-signal-degrade</b>                                                         | NA                                      | 15.1F5                                                             |
|                | <b>q-threshold-signal-degrade-clear</b>                                                   | NA                                      | 15.1F5                                                             |
| tca            | <b>odu-tca-bbe (enable-tca<br/>  no-enable-tca  <br/>threshold  <br/>threshold-24hrs)</b> | 13.3R3                                  | 15.1F5                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement | Options                                                                   | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-----------|---------------------------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
|           | odu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 13.3R3                                  | 15.1F5                                                             |
|           | odu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 13.3R3                                  | 15.1F5                                                             |
|           | odu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 13.3R3                                  | 15.1F5                                                             |
|           | odu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 13.3R3                                  | 15.1F5                                                             |
|           | odu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 13.3R3                                  | 15.1F5                                                             |
|           | odu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 13.3R3                                  | 15.1F5                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement | Options                                                                    | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-----------|----------------------------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
|           | otu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 13.3R3                                  | 15.1F5                                                             |
|           | otu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 13.3R3                                  | 15.1F5                                                             |
|           | otu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)      | 13.3R3                                  | 15.1F5                                                             |
|           | otu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)   | 13.3R3                                  | 15.1F5                                                             |
|           | otu-tca-fec-ber (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 13.3R3                                  | 15.1F5                                                             |
|           | otu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 13.3R3                                  | 15.1F5                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement                         | Options                                                                   | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-----------------------------------|---------------------------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
|                                   | otu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 13.3R3                                  | 15.1F5                                                             |
|                                   | otu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 13.3R3                                  | 15.1F5                                                             |
|                                   | otu-tca-uas-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 13.3R3                                  | 15.1F5                                                             |
| <i>transport-monitoring</i>       | -                                                                         | NA                                      | 15.1F5                                                             |
| <i>trigger-trigger-identifier</i> | oc-lof (hold-time (down   up)   ignore)                                   | 13.3R3                                  | 15.1F5                                                             |
|                                   | oc-lom (hold-time (down   up)   ignore)                                   | 13.3R3                                  | 15.1F5                                                             |
|                                   | oc-los (hold-time (down   up)   ignore)                                   | 13.3R3                                  | 15.1F5                                                             |
|                                   | oc-tsf (hold-time (down   up)   ignore)                                   | NA                                      | 15.1F5                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement | Options                                                    | MIC6-100G-CFP2<br>(MX2010 / MX2020) | MIC3-100G-DWDM<br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-----------|------------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
|           | <b>oc-wavelength-lock (hold-time (down   up)   ignore)</b> | NA                                  | 15.1F5                                                         |
|           | <b>odu-ais (hold-time (down   up)   ignore)</b>            | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-bdi (hold-time (down   up)   ignore)</b>            | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-bei (hold-time (down   up)   ignore)</b>            | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-iae (hold-time (down   up)   ignore)</b>            | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-lck (hold-time (down   up)   ignore)</b>            | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-oci (hold-time (down   up)   ignore)</b>            | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-sd (hold-time (down   up)   ignore)</b>             | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-tca-es</b>                                          | 13.3R3                              | NA                                                             |
|           | <b>odu-tca-ses</b>                                         | 13.3R3                              | NA                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement | Options                                                 | MIC6-100G-CFP2<br>(MX2010 / MX2020) | MIC3-100G-DWDM<br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-----------|---------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
|           | <b>odu-tca-uas</b>                                      | 13.3R3                              | NA                                                             |
|           | <b>odu-ttim</b>                                         | 13.3R3                              | 15.1F5                                                         |
|           | <b>opu-ptim (hold-time<br/>(down   up)   ignore)</b>    | 13.3R3                              | 15.1F5                                                         |
|           | <b>otu-ais (hold-time<br/>(down   up)   ignore)</b>     | 13.3R3                              | 15.1F5                                                         |
|           | <b>otu-bdi (hold-time<br/>(down   up)   ignore)</b>     | 13.3R3                              | 15.1F5                                                         |
|           | <b>otu-fec-deg (hold-time<br/>(down   up)   ignore)</b> | NA                                  | 15.1F5                                                         |
|           | <b>otu-fec-exe (hold-time<br/>(down   up)   ignore)</b> | NA                                  | 15.1F5                                                         |
|           | <b>otu-iae (hold-time<br/>(down   up)   ignore)</b>     | 13.3R3                              | 15.1F5                                                         |
|           | <b>otu-sd (hold-time (down<br/>  up)   ignore)</b>      | 13.3R3                              | 15.1F5                                                         |
|           | <b>odu-tca-es</b>                                       | 13.3R3                              | NA                                                             |
|           | <b>odu-tca-ses</b>                                      | 13.3R3                              | NA                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (Continued)

| Statement                 | Options                                                                                                      | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|---------------------------|--------------------------------------------------------------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
|                           | <b>odu-tca-uas</b>                                                                                           | 13.3R3                                  | NA                                                                 |
|                           | <b>otu-ttim (hold-time<br/>(down   up)   ignore)</b>                                                         | 13.3R3                                  | 15.1F5                                                             |
| <i>tti tti-identifier</i> | <b>odu-dapi <i>identifier</i></b>                                                                            | 13.3R3                                  | 15.1F5                                                             |
|                           | <b>odu-dapi-first-byte-nul  <br/>no-odu-dapi-first-byte-<br/>nul</b>                                         | NA                                      | 15.1F5                                                             |
|                           | <b>odu-expected-receive-<br/>dapi <i>identifier</i></b>                                                      | 13.3R3                                  | 15.1F5                                                             |
|                           | <b>odu-expected-receive-<br/>dapi-first-byte-nul   no-<br/>odu-expected-receive-<br/>dapi-first-byte-nul</b> | NA                                      | 15.1F5                                                             |
|                           | <b>odu-expected-receive-<br/>sapi <i>identifier</i></b>                                                      | 13.3R3                                  | 15.1F5                                                             |
|                           | <b>odu-sapi <i>identifier</i></b>                                                                            | 13.3R3                                  | 15.1F5                                                             |
|                           | <b>odu-sapi-first-byte-nul  <br/>no-odu-sapi-first-byte-<br/>nul</b>                                         | NA                                      | 15.1F5                                                             |
|                           | <b>otu-dapi <i>identifier</i></b>                                                                            | 13.3R3                                  | 15.1F5                                                             |

Table 96: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers (*Continued*)

| Statement | Options                                                                                                     | MIC6-100G-CFP2<br><br>(MX2010 / MX2020) | MIC3-100G-DWDM<br><br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-----------|-------------------------------------------------------------------------------------------------------------|-----------------------------------------|--------------------------------------------------------------------|
|           | <b>otu-dapi-first-byte-nul  <br/>no-otu-dapi-first-byte-<br/>nul</b>                                        | NA                                      | 15.1F5                                                             |
|           | <b>otu-expected- receive-<br/>dapi <i>identifier</i></b>                                                    | 13.3R3                                  | 15.1F5                                                             |
|           | <b>otu-expected-receive-<br/>dapi-first-byte-nul  no-<br/>odu-expected-receive-<br/>dapi-first-byte-nul</b> | NA                                      | 15.1F5                                                             |
|           | <b>otu-expected- receive-<br/>sapi <i>identifier</i></b>                                                    | 13.3R3                                  | 15.1F5                                                             |
|           | <b>otu-expected-receive-<br/>sapi-first-byte-nul  no-<br/>odu-expected-receive-<br/>dapi-first-byte-nul</b> | NA                                      | 15.1F5                                                             |
|           | <b>otu-sapi <i>identifier</i></b>                                                                           | 13.3R3                                  | 15.1F5                                                             |
|           | <b>otu-sapi-first-byte-nul</b>                                                                              | NA                                      | 15.1F5                                                             |

**SEE ALSO**| [Configuring OTN | 556](#)

## Supported OTN Options on ACX6360 Routers

Table 97 on page 515 lists the statements that are supported on ACX6360 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 97: Statements Supported on ACX6360 Routers**

| Statement                                            | Options                                        | Junos Version |
|------------------------------------------------------|------------------------------------------------|---------------|
| bytes (otn-options)                                  | transmit-payload- type <i>value</i>            | 18.3R1        |
| insert- odu-lck                                      | -                                              | 18.3R1        |
| insert- odu-oci                                      | -                                              | 18.3R1        |
| is-ma   no-is-ma                                     | -                                              | 18.3R1        |
| line-loopback   no-line-loopback                     | -                                              | 18.3R1        |
| local-loopback   no-local-loopback                   | -                                              | 18.3R1        |
| odu-ttim-action- enable   no-odu-ttim-action- enable | -                                              | 18.3R1        |
| otu-ttim-action- enable   no-otu-ttim-action- enable | -                                              | 18.3R1        |
| prbs   no-prbs                                       | -                                              | 18.3R1        |
| preemptive-fast-reroute                              | backward-frr -enable   no-backward-frr- enable | 18.3R1        |

Table 97: Statements Supported on ACX6360 Routers (Continued)

| Statement  | Options                                                                               | Junos Version |
|------------|---------------------------------------------------------------------------------------|---------------|
|            | signal- degrade- monitor- enable<br>  no-signal- degrade- monitor-<br>enable          | 18.3R1        |
|            | odu-backward- frr-enable   no-odu-<br>backward- frr-enable                            | NA            |
|            | odu-signal-degrade- monitor-<br>enable   no-odu-signal-degrade-<br>monitor-enable     | NA            |
| <b>tca</b> | <b>odu-tca-bbe (enable-tca   no-enable-tca<br/>  threshold   threshold-24hrs)</b>     | 18.3R1        |
|            | <b>odu-tca-bbe-fe (enable-tca   no-enable-<br/>tca   threshold   threshold-24hrs)</b> | 18.3R1        |
|            | <b>odu-tca-es (enable-tca   no-enable-tca  <br/>threshold   threshold-24hrs)</b>      | 18.3R1        |
|            | <b>odu-tca-es-fe (enable-tca   no-enable-<br/>tca   threshold   threshold-24hrs)</b>  | 18.3R1        |
|            | <b>odu-tca-ses (enable-tca   no-enable-tca<br/>  threshold   threshold-24hrs)</b>     | 18.3R1        |
|            | <b>odu-tca-ses-fe (enable-tca   no-enable-<br/>tca   threshold   threshold-24hrs)</b> | 18.3R1        |
|            | <b>odu-tca-uas (enable-tca   no-enable-tca<br/>  threshold   threshold-24hrs)</b>     | 18.3R1        |

Table 97: Statements Supported on ACX6360 Routers (Continued)

| Statement                                      | Options                                                                   | Junos Version |
|------------------------------------------------|---------------------------------------------------------------------------|---------------|
|                                                | otu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1        |
|                                                | otu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1        |
|                                                | otu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1        |
|                                                | otu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1        |
|                                                | otu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1        |
|                                                | otu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1        |
|                                                | otu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1        |
|                                                | otu-tca-uas-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1        |
| <code>trigger <i>trigger-identifier</i></code> | oc-lof (hold-time (down   up)   ignore)                                   | 18.3R1        |
|                                                | oc-lom (hold-time (down   up)   ignore)                                   | 18.3R1        |
|                                                | oc-tsfc (hold-time (down   up)   ignore)                                  | 18.3R1        |
|                                                | odu-ais (hold-time (down   up)   ignore)                                  | 18.3R1        |

Table 97: Statements Supported on ACX6360 Routers (Continued)

| Statement                 | Options                                          | Junos Version |
|---------------------------|--------------------------------------------------|---------------|
|                           | <b>odu-bdi (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>odu-bei (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>odu-iae (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>odu-lck (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>odu-oci (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>odu-sd (hold-time (down   up)   ignore)</b>   | 18.3R1        |
|                           | <b>odu-ttim</b>                                  | 18.3R1        |
|                           | <b>opu-ptim (hold-time (down   up)   ignore)</b> | 18.3R1        |
|                           | <b>otu-ais (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>otu-bdi (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>otu-iae (hold-time (down   up)   ignore)</b>  | 18.3R1        |
|                           | <b>otu-sd (hold-time (down   up)   ignore)</b>   | 18.3R1        |
|                           | <b>otu-ttim (hold-time (down   up)   ignore)</b> | 18.3R1        |
| <i>tti tti-identifier</i> | <b>odu-dapi identifier</b>                       | 18.3R1        |

Table 97: Statements Supported on ACX6360 Routers (Continued)

| Statement | Options                                             | Junos Version |
|-----------|-----------------------------------------------------|---------------|
|           | <b>odu-expected- receive-dapi <i>identifier</i></b> | 18.3R1        |
|           | <b>odu-expected- receive-sapi <i>identifier</i></b> | 18.3R1        |
|           | <b>odu-sapi <i>identifier</i></b>                   | 18.3R1        |
|           | <b>otu-dapi <i>identifier</i></b>                   | 18.3R1        |
|           | <b>otu-expected- receive-dapi <i>identifier</i></b> | 18.3R1        |
|           | <b>otu-expected- receive-sapi <i>identifier</i></b> | 18.3R1        |
|           | <b>otu-sapi <i>identifier</i></b>                   | 18.3R1        |

## SEE ALSO

| [Configuring OTN | 556](#)

## Supported OTN Options on ACX5448-D Routers

Table 98 on page 519 lists the statements that are supported on ACX5448-D routers at the [edit interfaces *interface-name* otn-options] hierarchy level.

Table 98: Statements Supported on ACX5448-D Routers

| Statement                        | Options                                    | Junos OS |
|----------------------------------|--------------------------------------------|----------|
| <code>bytes (otn-options)</code> | <b>transmit-payload- type <i>value</i></b> | 19.2R1   |
| <b>insert- odu-lck</b>           | -                                          | 19.2R1   |

Table 98: Statements Supported on ACX5448-D Routers (Continued)

| Statement                                                  | Options                                                                      | Junos OS |
|------------------------------------------------------------|------------------------------------------------------------------------------|----------|
| <b>insert- odu-oci</b>                                     | -                                                                            | 19.2R1   |
| is-ma   <b>no-is-ma</b>                                    | -                                                                            | 19.2R1   |
| line-loopback   <b>no-line-loopback</b>                    | -                                                                            | 19.2R1   |
| local-loopback   <b>no-local-loopback</b>                  | -                                                                            | 19.2R1   |
| odu-ttim-action- enable   <b>no-odu-ttim-action-enable</b> | -                                                                            | 19.2R1   |
| otu-ttim-action- enable   <b>no-otu-ttim-action-enable</b> | -                                                                            | 19.2R1   |
| prbs   <b>no-prbs</b>                                      | -                                                                            | 19.2R1   |
| preemptive-fast-reroute                                    | backward-frr -enable   <b>no-backward-frr- enable</b>                        | 19.2R1   |
|                                                            | signal- degrade- monitor- enable   <b>no-signal- degrade- monitor-enable</b> | 19.2R1   |
|                                                            | odu-backward- frr-enable   <b>no-odu-backward- frr-enable</b>                | NA       |

Table 98: Statements Supported on ACX5448-D Routers (Continued)

| Statement  | Options                                                                          | Junos OS |
|------------|----------------------------------------------------------------------------------|----------|
|            | odu-signal-degrade-monitor-enable   <b>no</b> -odu-signal-degrade-monitor-enable | NA       |
| <b>tca</b> | <b>odu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 19.2R1   |
|            | <b>odu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 19.2R1   |
|            | <b>odu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>     | 19.2R1   |
|            | <b>odu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>  | 19.2R1   |
|            | <b>odu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 19.2R1   |
|            | <b>odu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 19.2R1   |
|            | <b>odu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 19.2R1   |
|            | <b>otu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 19.2R1   |
|            | <b>otu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 19.2R1   |

Table 98: Statements Supported on ACX5448-D Routers (*Continued*)

| Statement                                      | Options                                                                   | Junos OS |
|------------------------------------------------|---------------------------------------------------------------------------|----------|
|                                                | otu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 19.2R1   |
|                                                | otu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 19.2R1   |
|                                                | otu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 19.2R1   |
|                                                | otu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 19.2R1   |
|                                                | otu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 19.2R1   |
|                                                | otu-tca-uas-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 19.2R1   |
| <code>trigger <i>trigger-identifier</i></code> | oc-lof (hold-time (down   up)   ignore)                                   | 19.2R1   |
|                                                | oc-lom (hold-time (down   up)   ignore)                                   | 19.2R1   |
|                                                | oc-tsf (hold-time (down   up)   ignore)                                   | 19.2R1   |
|                                                | odu-ais (hold-time (down   up)   ignore)                                  | 19.2R1   |
|                                                | odu-bdi (hold-time (down   up)   ignore)                                  | 19.2R1   |
|                                                | odu-bei (hold-time (down   up)   ignore)                                  | 19.2R1   |

Table 98: Statements Supported on ACX5448-D Routers (Continued)

| Statement                 | Options                                      | Junos OS |
|---------------------------|----------------------------------------------|----------|
|                           | odu-iae (hold-time (down   up)   ignore)     | 19.2R1   |
|                           | odu-lck (hold-time (down   up)   ignore)     | 19.2R1   |
|                           | odu-oci (hold-time (down   up)   ignore)     | 19.2R1   |
|                           | odu-sd (hold-time (down   up)   ignore)      | 19.2R1   |
|                           | odu-ttim                                     | 19.2R1   |
|                           | opu-ptim (hold-time (down   up)   ignore)    | 19.2R1   |
|                           | otu-ais (hold-time (down   up)   ignore)     | 19.2R1   |
|                           | otu-bdi (hold-time (down   up)   ignore)     | 19.2R1   |
|                           | otu-iae (hold-time (down   up)   ignore)     | 19.2R1   |
|                           | otu-sd (hold-time (down   up)   ignore)      | 19.2R1   |
|                           | otu-ttim (hold-time (down   up)   ignore)    | 19.2R1   |
| <i>tti tti-identifier</i> | odu-dapi <i>identifier</i>                   | 19.2R1   |
|                           | odu-expected- receive-dapi <i>identifier</i> | 19.2R1   |
|                           | odu-expected- receive-sapi <i>identifier</i> | 19.2R1   |

Table 98: Statements Supported on ACX5448-D Routers (*Continued*)

| Statement | Options                                             | Junos OS |
|-----------|-----------------------------------------------------|----------|
|           | <b>odu-sapi <i>identifier</i></b>                   | 19.2R11  |
|           | <b>otu-dapi <i>identifier</i></b>                   | 19.2R1   |
|           | <b>otu-expected- receive-dapi <i>identifier</i></b> | 19.2R1   |
|           | <b>otu-expected- receive-sapi <i>identifier</i></b> | 19.2R1   |
|           | <b>otu-sapi <i>identifier</i></b>                   | 19.2R1   |

## SEE ALSO

[Configuring OTN | 556](#)

## Supported OTN Options on PTX10008 and PTX10016 Series Routers

[Table 99 on page 524](#) lists the statements that are supported on the PTX10K-LC1104 line card on PTX10008 and PTX10016 routers at the `[edit interfaces interface-name otn-options]` hierarchy level.

Table 99: Statements Supported on PTX10K-LC1104 line cards

| Statement                        | Options                                                                            | Release | Interfaces Supported<br>(ot/et) |
|----------------------------------|------------------------------------------------------------------------------------|---------|---------------------------------|
| <code>bytes (otn-options)</code> | <b>transmit-payload- type <i>value</i></b>                                         | 18.3R1  | et                              |
| <code>fec</code>                 | <b>(<i>efec</i> / <i>gfec</i> / <i>gfec-sdfec</i> / <i>none</i> / <i>ufec</i>)</b> | 18.3R1  | ot                              |

Table 99: Statements Supported on PTX10K-LC1104 line cards (Continued)

| Statement                                    | Options                                              | Release | Interfaces Supported<br>(ot/et) |
|----------------------------------------------|------------------------------------------------------|---------|---------------------------------|
| <b>insert- odu-lck</b>                       | -                                                    | 18.3R1  | et                              |
| <b>insert- odu-oci</b>                       | -                                                    | 18.3R1  | et                              |
| is-ma   <b>no-is-ma</b>                      | -                                                    | 18.3R1  | et                              |
| laser-enable  <br><b>no-laser-enable</b>     | -                                                    | 18.3R1  | ot                              |
| line-loopback  <br><b>no-line-loopback</b>   | -                                                    | 18.3R1  | ot                              |
| local-loopback  <br><b>no-local-loopback</b> | -                                                    | 18.3R1  | et                              |
| <i>modulation-format</i>                     | <b>qpsk   8qam   16qam</b>                           | 18.3R1  | ot                              |
| odu-delay-management                         | bypass   <b>no-bypass</b>                            | 18.3R1  | et                              |
|                                              | monitor- end-point  <br><b>no-monitor-end-point</b>  | 18.3R1  |                                 |
|                                              | number- of-frames<br><i>value</i>                    | 18.3R1  |                                 |
|                                              | <b>no-start- measurement</b>  <br>start- measurement | 18.3R1  |                                 |

Table 99: Statements Supported on PTX10K-LC1104 line cards (Continued)

| Statement                                                       | Options                                                                                 | Release | Interfaces Supported<br>(ot/et) |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------|---------|---------------------------------|
| <b>odu-signal-degrade</b>                                       | ber-threshold -clear <i>value</i>                                                       | 18.3R1  | et                              |
|                                                                 | ber-threshold -signal-degrade <i>value</i>                                              | 18.3R1  |                                 |
|                                                                 | interval <i>value</i>                                                                   | 18.3R1  |                                 |
| odu-ttim-action- enable  <br><b>no</b> -odu-ttim-action- enable | -                                                                                       | 18.3R1  | et                              |
| otu-ttim-action- enable  <br><b>no</b> -otu-ttim-action- enable | -                                                                                       | 18.3R1  | et                              |
| prbs   <b>no</b> -prbs                                          | -                                                                                       | 18.3R1  | ot                              |
| preemptive-fast- reroute                                        | backward-frr -enable  <br><b>no</b> -backward-frr-<br>enable                            | 18.3R1  | ot and et                       |
|                                                                 | signal- degrade-<br>monitor- enable   <b>no</b> -<br>signal- degrade-<br>monitor-enable | 18.3R1  |                                 |

Table 99: Statements Supported on PTX10K-LC1104 line cards (Continued)

| Statement      | Options                                                                                      | Release | Interfaces Supported<br>(ot/et) |
|----------------|----------------------------------------------------------------------------------------------|---------|---------------------------------|
|                | odu-backward- frr-<br>enable   <b>no</b> -odu-<br>backward- frr-enable                       | 18.3R1  |                                 |
|                | odu-signal-degrade-<br>monitor-enable   <b>no</b> -<br>odu-signal-degrade-<br>monitor-enable | 18.3R1  |                                 |
| rate           | fixed-stuff-bytes   <b>no</b> -<br>fixed-stuff-bytes                                         | 18.3R1  | et                              |
|                | <b>oc192</b>                                                                                 | 18.3R1  |                                 |
|                | otu4                                                                                         | 18.3R1  |                                 |
|                | pass-through   <b>no</b> -pass-<br>through                                                   | 18.3R1  |                                 |
| signal-degrade | ber-threshold-clear<br><b>value</b>                                                          | 18.3R1  | ot and et                       |
|                | ber-threshold-signal-<br>degrade <b>value</b>                                                | 18.3R1  |                                 |
|                | interval <b>value</b>                                                                        | 18.3R1  |                                 |

Table 99: Statements Supported on PTX10K-LC1104 line cards *(Continued)*

| Statement  | Options                                                                          | Release | Interfaces Supported<br><br>(ot/et) |
|------------|----------------------------------------------------------------------------------|---------|-------------------------------------|
| <b>tca</b> | <b>odu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 18.3R1  | ot and et                           |
|            | <b>odu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 18.3R1  |                                     |
|            | <b>odu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>     | 18.3R1  |                                     |
|            | <b>odu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>  | 18.3R1  |                                     |
|            | <b>odu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 18.3R1  |                                     |
|            | <b>odu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 18.3R1  |                                     |
|            | <b>odu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>    | 18.3R1  |                                     |

Table 99: Statements Supported on PTX10K-LC1104 line cards (Continued)

| Statement | Options                                                                    | Release | Interfaces Supported<br><br>(ot/et) |
|-----------|----------------------------------------------------------------------------|---------|-------------------------------------|
|           | otu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                                     |
|           | otu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1  |                                     |
|           | otu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)      | 18.3R1  |                                     |
|           | otu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)   | 18.3R1  |                                     |
|           | otu-tca-fec-ber (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                                     |
|           | otu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                                     |
|           | otu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1  |                                     |

Table 99: Statements Supported on PTX10K-LC1104 line cards (*Continued*)

| Statement                         | Options                                                                   | Release | Interfaces Supported<br><br>(ot/et) |
|-----------------------------------|---------------------------------------------------------------------------|---------|-------------------------------------|
|                                   | otu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                                     |
|                                   | otu-tca-uas-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                                     |
| <i>transport-monitoring</i>       | -                                                                         | 18.3R1  | et                                  |
| <i>trigger trigger-identifier</i> | oc-lof (hold-time (down   up)   ignore)                                   | 18.3R1  | ot and et                           |
|                                   | oc-lom (hold-time (down   up)   ignore)                                   | 18.3R1  |                                     |
|                                   | oc-los (hold-time (down   up)   ignore)                                   | 18.3R1  |                                     |
|                                   | oc-tsf (hold-time (down   up)   ignore)                                   | 18.3R1  |                                     |
|                                   | oc-wavelength-lock (hold-time (down   up)   ignore)                       | 18.3R1  |                                     |
|                                   | odu-ais (hold-time (down   up)   ignore)                                  | 18.3R1  |                                     |

Table 99: Statements Supported on PTX10K-LC1104 line cards *(Continued)*

| Statement | Options                                   | Release | Interfaces Supported<br><br>(ot/et) |
|-----------|-------------------------------------------|---------|-------------------------------------|
|           | odu-bdi (hold-time (down   up)   ignore)  | 18.3R1  |                                     |
|           | odu-bei (hold-time (down   up)   ignore)  | 18.3R1  |                                     |
|           | odu-iae (hold-time (down   up)   ignore)  | 18.3R1  |                                     |
|           | odu-lck (hold-time (down   up)   ignore)  | 18.3R1  |                                     |
|           | odu-oci (hold-time (down   up)   ignore)  | 18.3R1  |                                     |
|           | odu-sd (hold-time (down   up)   ignore)   | 18.3R1  |                                     |
|           | odu-ttim                                  | 18.3R1  |                                     |
|           | opu-ptim (hold-time (down   up)   ignore) | 18.3R1  |                                     |
|           | otu-ais (hold-time (down   up)   ignore)  | 18.3R1  |                                     |
|           | otu-bdi (hold-time (down   up)   ignore)  | 18.3R1  |                                     |

Table 99: Statements Supported on PTX10K-LC1104 line cards (*Continued*)

| Statement                 | Options                                      | Release | Interfaces Supported<br>(ot/et) |
|---------------------------|----------------------------------------------|---------|---------------------------------|
|                           | otu-fec-deg (hold-time (down   up)   ignore) | 18.3R1  |                                 |
|                           | otu-fec-exe (hold-time (down   up)   ignore) | 18.3R1  |                                 |
|                           | otu-iae (hold-time (down   up)   ignore)     | 18.3R1  |                                 |
|                           | otu-sd (hold-time (down   up)   ignore)      | 18.3R1  |                                 |
|                           | otu-ttim (hold-time (down   up)   ignore)    | 18.3R1  |                                 |
| <i>tti tti-identifier</i> | odu-dapi <i>identifier</i>                   | 18.3R1  | et                              |
|                           | odu-expected- receive-dapi <i>identifier</i> | 18.3R1  |                                 |
|                           | odu-expected- receive-sapi <i>identifier</i> | 18.3R1  |                                 |
|                           | odu-sapi <i>identifier</i>                   | 18.3R1  |                                 |
|                           | otu-dapi <i>identifier</i>                   | 18.3R1  |                                 |

Table 99: Statements Supported on PTX10K-LC1104 line cards (Continued)

| Statement | Options                                                  | Release | Interfaces Supported<br>(ot/et) |
|-----------|----------------------------------------------------------|---------|---------------------------------|
|           | <b>otu-expected- receive-<br/>dapi <i>identifier</i></b> | 18.3R1  |                                 |
|           | <b>otu-expected- receive-<br/>sapi <i>identifier</i></b> | 18.3R1  |                                 |
|           | <b>otu-sapi <i>identifier</i></b>                        | 18.3R1  |                                 |

### Supported Optics Options on ACX6360 and ACX5448-D Routers

Table 100 on page 533 lists the statements that are supported on ACX6360 and ACX5448-D routers at the [edit interfaces *interface-name* optics-options] hierarchy level.

Table 100: Statements Supported on ACX6360 and ACX5448-D Routers

| Statement                         | Options                | Release           | Interfaces Supported |
|-----------------------------------|------------------------|-------------------|----------------------|
| <i>fec</i>                        | <b>sdfec   sdfec15</b> | 18.3R1, 19.2R1-S1 | ot                   |
| high-polarization                 | -                      | 18.2R1, 19.2R1-S1 | ot                   |
| laser-enable  <br>no-laser-enable | -                      | 18.2R1, 19.2R1-S1 | ot                   |
| los-alarm-<br>threshold           | -                      | 18.2R1, 19.2R1-S1 | ot                   |
| los-warn-threshold                | -                      | 18.2R1, 19.2R1-S1 | ot                   |

Table 100: Statements Supported on ACX6360 and ACX5448-D Routers (Continued)

| Statement                | Options                                                                                             | Release           | Interfaces Supported |
|--------------------------|-----------------------------------------------------------------------------------------------------|-------------------|----------------------|
| <i>modulation-format</i> | <b>(16qam   8qam   qpsk)</b>                                                                        | 18.2R1, 19.2R1-S1 | ot                   |
| signal-degrade           | ber-threshold-clear <b>value</b>                                                                    | 18.3R1, 19.2R1-S1 | ot                   |
|                          | ber-threshold-signal-degrade <b>value</b>                                                           | 18.3R1, 19.2R1-S1 |                      |
|                          | interval <b>value</b>                                                                               | 18.3R1, 19.2R1-S1 |                      |
|                          | <b>q-threshold-signal-degrade</b>                                                                   | 18.3R1, 19.2R1-S1 |                      |
|                          | <b>q-threshold-signal-degrade-clear</b>                                                             | 18.3R1, 19.2R1-S1 |                      |
| tca                      | <b>carrier-frequency-offset-high-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b> | 18.2R1, 19.2R1-S1 | ot                   |
|                          | <b>carrier-frequency-offset-low-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>  | 18.2R1, 19.2R1-S1 |                      |
|                          | <b>fec-ber (enable-tca   no-enable-tca   threshold   threshold-24hrs)</b>                           | 18.2R1, 19.2R1-S1 |                      |

Table 100: Statements Supported on ACX6360 and ACX5448-D Routers *(Continued)*

| Statement | Options                                                                                   | Release           | Interfaces Supported |
|-----------|-------------------------------------------------------------------------------------------|-------------------|----------------------|
|           | fec-corrected-errors-high-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.2R1, 19.2R1-S1 |                      |
|           | fec-uncorrected-words-high-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 |                      |
|           | laser-frequency-error-high-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 |                      |
|           | laser-frequency-error-low-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.2R1, 19.2R1-S1 |                      |
|           | pam-histogram-high-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)         | 18.2R1, 19.2R1-S1 |                      |
|           | residual-isi-high-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)          | 18.2R1, 19.2R1-S1 |                      |

Table 100: Statements Supported on ACX6360 and ACX5448-D Routers *(Continued)*

| Statement | Options                                                                            | Release           | Interfaces Supported |
|-----------|------------------------------------------------------------------------------------|-------------------|----------------------|
|           | residual-isi-low-tca<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 |                      |
|           | rx-power-high-tca<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.2R1, 19.2R1-S1 |                      |
|           | rx-power-low-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)        | 18.2R1, 19.2R1-S1 |                      |
|           | snr-low-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)             | 18.2R1, 19.2R1-S1 |                      |
|           | tec-current-high-tca<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 |                      |
|           | tec-current-low-tca<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.2R1, 19.2R1-S1 |                      |

Table 100: Statements Supported on ACX6360 and ACX5448-D Routers (*Continued*)

| Statement         | Options                                                                            | Release           | Interfaces Supported |
|-------------------|------------------------------------------------------------------------------------|-------------------|----------------------|
|                   | temperature-high-tca<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 |                      |
|                   | temperature-low-tca<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.2R1, 19.2R1-S1 |                      |
|                   | tx-power-high-tca<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.2R1, 19.2R1-S1 |                      |
|                   | tx-power-low-tca (enable-tca   no-enable-tca   threshold   threshold-24hrs)        | 18.2R1, 19.2R1-S1 |                      |
| <i>tx-power</i>   | dbm                                                                                | 18.2R1, 19.2R1-S1 | ot                   |
| <i>wavelength</i> | nm                                                                                 | 18.2R1, 19.2R1-S1 | ot                   |

**SEE ALSO**

| [Configuring OTN | 556](#)

**Supported Optics Options on PTX10008 and PTX10016 Series Routers**

Table 101 on page 538 lists the statements that are supported on PTX10008 and PTX10016 Series routers at the [edit interfaces *interface-name* optics-options] hierarchy level.

**Table 101: Statements Supported on PTX10008 and PTX10016 Series Routers**

| Statement                             | Options                                                    | Release | Interfaces Supported |
|---------------------------------------|------------------------------------------------------------|---------|----------------------|
| alarm low-light-alarm                 | link-down   syslog                                         | 18.3R1  | ot                   |
| tca carrier-frequency-offset-high-tca | (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  | ot                   |
| tx-power                              | dbm                                                        | 18.3R1  | ot                   |
| warning low-light-warning             | link-down   syslog                                         | 18.3R1  | ot                   |
| laser-enable   no-laser-enable        | -                                                          | 18.3R1  | ot                   |
| line-loopback   no-line-loopback      | -                                                          | 18.3R1  | ot                   |
| prbs   no-prbs                        | -                                                          | 18.3R1  | ot                   |
| signal-degrade                        | ber-threshold-clear <i>value</i>                           | 18.3R1  | ot                   |
|                                       | ber-threshold-signal-degrade <i>value</i>                  | 18.3R1  |                      |
|                                       | interval <i>value</i>                                      | 18.3R1  |                      |

Table 101: Statements Supported on PTX10008 and PTX10016 Series Routers *(Continued)*

| Statement | Options                                                                   | Release | Interfaces Supported |
|-----------|---------------------------------------------------------------------------|---------|----------------------|
| tca       | odu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  | ot                   |
|           | odu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                      |
|           | odu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                      |
|           | odu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1  |                      |
|           | odu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | odu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                      |
|           | odu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                      |

Table 101: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

| Statement | Options                                                                   | Release | Interfaces Supported |
|-----------|---------------------------------------------------------------------------|---------|----------------------|
|           | otu-tca-bbe (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | otu-tca-bbe-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                      |
|           | otu-tca-es (enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                      |
|           | otu-tca-es-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1  |                      |
|           | otu-tca-ses (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | otu-tca-ses-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                      |
|           | otu-tca-uas (enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | otu-tca-uas-fe (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                      |

Table 101: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

| Statement                                      | Options                                                          | Release | Interfaces Supported |
|------------------------------------------------|------------------------------------------------------------------|---------|----------------------|
| <code>trigger <i>trigger-identifier</i></code> | <code>oc-lof (hold-time (down   up)   ignore)</code>             | 18.3R1  | ot                   |
|                                                | <code>oc-lom (hold-time (down   up)   ignore)</code>             | 18.3R1  |                      |
|                                                | <code>oc-los (hold-time (down   up)   ignore)</code>             | 18.3R1  |                      |
|                                                | <code>oc-tsf (hold-time (down   up)   ignore)</code>             | 18.3R1  |                      |
|                                                | <code>oc-wavelength-lock (hold-time (down   up)   ignore)</code> | 18.3R1  |                      |
|                                                | <code>odu-ais (hold-time (down   up)   ignore)</code>            | 18.3R1  |                      |
|                                                | <code>odu-bdi (hold-time (down   up)   ignore)</code>            | 18.3R1  |                      |
|                                                | <code>odu-bei (hold-time (down   up)   ignore)</code>            | 18.3R1  |                      |
|                                                | <code>odu-iae (hold-time (down   up)   ignore)</code>            | 18.3R1  |                      |
|                                                | <code>odu-lck (hold-time (down   up)   ignore)</code>            | 18.3R1  |                      |

Table 101: Statements Supported on PTX10008 and PTX10016 Series Routers (Continued)

| Statement | Options                                      | Release | Interfaces Supported |
|-----------|----------------------------------------------|---------|----------------------|
|           | odu-oci (hold-time (down   up)   ignore)     | 18.3R1  |                      |
|           | odu-sd (hold-time (down   up)   ignore)      | 18.3R1  |                      |
|           | odu-ttim                                     | 18.3R1  |                      |
|           | opu-ptim (hold-time (down   up)   ignore)    | 18.3R1  |                      |
|           | otu-ais (hold-time (down   up)   ignore)     | 18.3R1  |                      |
|           | otu-bdi (hold-time (down   up)   ignore)     | 18.3R1  |                      |
|           | otu-fec-deg (hold-time (down   up)   ignore) | 18.3R1  |                      |
|           | otu-fec-exe (hold-time (down   up)   ignore) | 18.3R1  |                      |
|           | otu-iae (hold-time (down   up)   ignore)     | 18.3R1  |                      |
|           | otu-sd (hold-time (down   up)   ignore)      | 18.3R1  |                      |

Table 101: Statements Supported on PTX10008 and PTX10016 Series Routers (*Continued*)

| Statement | Options                                          | Release | Interfaces Supported |
|-----------|--------------------------------------------------|---------|----------------------|
|           | <b>otu-ttim (hold-time (down   up)   ignore)</b> | 18.3R1  |                      |

## RELATED DOCUMENTATION

[Configuring OTN | 556](#)

## Forward Error Correction (FEC) and Bit Error Rate (BER)

### IN THIS SECTION

- [Understanding Pre-FEC BER Monitoring and BER Thresholds | 543](#)
- [Supported Forward Error Correction Modes on MX Series Routers | 549](#)
- [Supported Forward Error Correction Modes on PTX Series Routers | 549](#)
- [Supported Forward Error Correction Modes on ACX6360 Router | 550](#)
- [Supported FEC Modes on ACX5448-D Router | 551](#)

OTN interfaces use pre-forward error correction (Pre-FEC) bit error rate (BER) for monitoring the condition of an OTN link. Use this topic to understand more about how OTN links are monitored and the supported FEC modes on devices.

### Understanding Pre-FEC BER Monitoring and BER Thresholds

Optical transport network (OTN) interfaces on PTX Series Packet Transport Routers support monitoring the condition of an OTN 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


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). Refer to *Supported Forward Error Correction Modes on ACX6360 Router* for more details.

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 7 on page 545](#).

After an alarm is raised, if the BER returns to a level below the threshold clear value (*ber-threshold-clear*), the alarm is cleared.

**Figure 7: Pre-FEC BER Monitoring**



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 OTN frames, notifying the remote interface of signal degradation. The remote interface can use the information to reroute traffic to a different interface. If you use pre-FEC BER monitoring together with backward fast reroute, then notification of signal degradation and rerouting of traffic occurs in less time than that required through a Layer 3 protocol.

Include the `signal-degrade-monitor-enable` and `backward-frr-enable` statements at the [edit interfaces *interface-name* otn-options preemptive-fast-reroute] hierarchy level to enable pre-FEC BER monitoring and backward fast reroute.

**NOTE:** When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the `signal-degrade-monitor-enable` and the `backward-frr-enable` statements.

You can also configure the pre-FEC BER thresholds that raise or clear a signal degrade alarm and the time interval for the thresholds. If the BER thresholds and interval are not configured, the default values are used.

When a received signal degrade alarm is active and backward fast reroute is enabled, a specific flag is inserted into the transmitted OTN overhead. The remote PIC at the opposite end of the link monitors the OTN 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 OTN 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<sup>2</sup>-factor measurements (also known as Q-factor). Q<sup>2</sup>-factor is expressed in units of decibels relative to a Q<sup>2</sup>-factor of zero (dBQ). Q<sup>2</sup>-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<sup>2</sup>-factor to BER to enter them in the CLI by using scientific notation. BER can be converted to Q<sup>2</sup>-factor by using the following equation:

$$Q^2\text{-factor} = 20 * \log_{10} (\text{sqrt}(2) * \text{erfcinv}(2 * \text{BER}))$$

**TIP:** To convert between Q<sup>2</sup>-factor and BER in a spreadsheet program, you can approximate the values by using the following formulas:

- To calculate Q<sup>2</sup>-factor:

$$= 20 * \text{LOG10} (-\text{NORMSINV}(\text{BER}))$$

- To calculate BER:

$$= 1 - \text{NORMSDIST}(10^{(0.05 * Q^2\text{-factor})})$$

Table 102 on page 547 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 102: Example—Signal Degrade and Clear Threshold Values at 1 dBQ**

| PIC                  | FEC Type           | FEC Limit              |        | Signal Degrade Threshold |        | Clear Threshold        |        |
|----------------------|--------------------|------------------------|--------|--------------------------|--------|------------------------|--------|
|                      |                    | Q <sup>2</sup> -Factor | BER    | Q <sup>2</sup> -Factor   | BER    | Q <sup>2</sup> -Factor | BER    |
| P1-PTX-2-100G-WDM    | SD-FEC             | 6.7 dBQ                | 1.5E-2 | 7.7 dBQ                  | 7.5E-3 | 8.7 dBQ                | 3.0E-3 |
| P2-100GE-OTN         | G.709 GFEC         | 11.5 dBQ               | 8.0E-5 | 12.5 dBQ                 | 1.1E-5 | 13.5 dBQ               | 1.0E-6 |
| P1-PTX-24-10G-W-SFPP | G.975.1 I.4 (UFEC) | 9.1 dBQ                | 2.2E-3 | 10.1 dBQ                 | 6.9E-4 | 11.1 dBQ               | 1.6E-4 |
|                      | G.975.1 I.7 (EFEC) | 9.6 dBQ                | 1.3E-3 | 10.6 dBQ                 | 3.6E-4 | 11.6 dBQ               | 7.5E-5 |
|                      | G.709 GFEC         | 11.5 dBQ               | 8.0E-5 | 12.5 dBQ                 | 1.1E-5 | 13.5 dBQ               | 1.0E-6 |

To adjust the signal degrade threshold, you must first decide on a new system margin target and then calculate the respective BER value (using the equation to convert from Q<sup>2</sup>-factor to BER). Table 103 on page 548 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).

**NOTE:** 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 103: Example—Signal Degrade and Clear Thresholds After Configuration**

| PIC                  | FEC Type           | FEC Limit              |        | Signal Degrade Threshold |        | Clear Threshold        |        |
|----------------------|--------------------|------------------------|--------|--------------------------|--------|------------------------|--------|
|                      |                    | Q <sup>2</sup> -Factor | BER    | Q <sup>2</sup> -Factor   | BER    | Q <sup>2</sup> -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 |

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.

**SEE ALSO**

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | [459](#)

**Supported Forward Error Correction Modes on MX Series Routers**

[Table 104 on page 549](#) lists the FEC modes that are supported on MX Series 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 line card:

**Table 104: FEC modes Supported on MX Series Routers**

| Line Card                      | FEC Mode                           | Port Speed                  | Junos Version |
|--------------------------------|------------------------------------|-----------------------------|---------------|
| <a href="#">MPC5E-40G10G</a>   | <i>(gfec / efec / none / ufec)</i> | 10G                         | 13.3          |
| <a href="#">MPC5E-100G10G</a>  | <i>(gfec / efec / none / ufec)</i> | 10G and 100G<br>(GFEC only) | 13.3          |
| <a href="#">MIC6-10G-OTN</a>   | <i>(gfec / efec / none / ufec)</i> | 10G                         | 13.3          |
| <a href="#">MIC6-100G-CFP2</a> | <i>(gfec / none )</i>              | 100G (GFEC only)            | 13.3          |
| <a href="#">MIC3-100G-DWDM</a> | <i>gfec / hgfec / sdfec</i>        | 100G                        | 15.1F5        |

**SEE ALSO**

[fec](#) | [709](#)

**Supported Forward Error Correction Modes on PTX Series Routers**

[Table 105 on page 550](#) lists the FEC modes that are supported on PTX Series routers at the `[edit interfaces interface-name otn-options]` hierarchy level.

Table 105: FEC Modes Supported on PTX Series Routers

| Line Card                            | FEC Mode                           | Port Speed       | Junos Version                                         |
|--------------------------------------|------------------------------------|------------------|-------------------------------------------------------|
| <a href="#">P1-PTX-24-10G-W-SFPP</a> | <i>(gfec / efec / none / ufec)</i> | 10G              | 12.1X48, 12.3, 13.2 (PTX5000)<br><br>13.2R2 (PTX3000) |
| <a href="#">P2-10G-40G-QSFPP</a>     | <i>(gfec / efec / none / ufec)</i> | 10G              | 14.1R2 (PTX5000)<br><br>15.1F6 (PTX3000)              |
| <a href="#">P2-100GE-OTN</a>         | <i>(gfec / none )</i>              | 100G (GFEC only) | 14.1                                                  |
| <a href="#">P1-PTX-2-100G-WDM</a>    | <i>(gfec-sdfec)</i>                | 100G             | 13.2 (PTX5000)<br><br>13.3 (PTX3000)                  |
| <a href="#">PTX-5-100G-WDM</a>       | <i>gfec / sdfec</i>                | 100G             | 15.1F6                                                |

## SEE ALSO

| [fec](#) | [709](#)

## Supported Forward Error Correction Modes on ACX6360 Router

[Table 106 on page 550](#) lists the FEC modes that are supported on ACX6360 routers at the `[edit interfaces interface-name optics-options]` hierarchy level.

Table 106: FEC modes Supported on ACX6360 Routers

| FEC Mode     | Modulation Format | Port Speed | Junos Version |
|--------------|-------------------|------------|---------------|
| <i>sdfec</i> | QPSK              | 100G       | 18.3R1        |

Table 106: FEC modes Supported on ACX6360 Routers (*Continued*)

| FEC Mode       | Modulation Format | Port Speed | Junos Version |
|----------------|-------------------|------------|---------------|
| <i>sdfec15</i> | QPSK              | 100G       | 18.3R1        |
| <i>sdfec15</i> | 8-QAM             | 200G       | 18.3R1        |
| <i>sdfec15</i> | 16-QAM            | 200G       | 18.3R1        |

## SEE ALSO

| [fec](#) | [709](#)

## Supported FEC Modes on ACX5448-D Router

Table 107 on page 551 lists the forward error correction (FEC) modes that are supported on ACX5448-D routers. You can configure the FEC modes at the [edit interfaces *interface-name* optics-options] hierarchy level.

Table 107: FEC Modes Supported on ACX5448-D Routers

| FEC Mode       | Modulation Format | Port Speed | Junos OS  |
|----------------|-------------------|------------|-----------|
| <b>sdfec</b>   | QPSK              | 100 Gbps   | 19.2R1-S1 |
| <b>hgfec</b>   | QPSK              | 100 Gbps   | 19.2R1-S1 |
| <b>sdfec15</b> | QPSK              | 100 Gbps   | 19.2R1-S1 |
| <b>sdfec15</b> | 8-QAM             | 200 Gbps   | 19.2R1-S1 |
| <b>sdfec15</b> | 16-QAM            | 200 Gbps   | 19.2R1-S1 |

**SEE ALSO**

| [fec](#) | [709](#)

**Release History Table**

| 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). |

**RELATED DOCUMENTATION**

| [Understanding Optical Transport Network \(OTN\)](#) | [458](#)

| [Configuring OTN](#) | [556](#)

## Dense Wavelength Division Multiplexing (DWDM) Interface Wavelength

**IN THIS SECTION**

- [Ethernet DWDM Interface Wavelength Overview](#) | [552](#)
- [Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength](#) | [553](#)

Use this topic to understand about the dense wavelength-division multiplexing (dwdm) wavelength parameter and how to configure the wavelength for 10-Gigabit and 100-Gigabit Ethernet Interfaces.

### Ethernet DWDM Interface Wavelength Overview

Dense wavelength-division multiplexing (DWDM) interfaces are supported on 10-Gigabit Ethernet DWDM PICs, MICs, and MPCs; the 10-Gigabit Ethernet LAN/WAN OTN PIC; and the 100-Gigabit Ethernet DWDM OTN PIC. When a tunable optic transceiver is available, you can configure the DWDM interfaces with full C-band International Telecommunication Union (ITU)-Grid tunable optics, as defined in the following specifications:

- *Intel TXN13600 Optical Transceiver I2C Interface and Customer EEPROM Preliminary Specification*, July 2004.

- *I2C Reference Document for 300-Pin MSA 10G and 40G Transponder*, Edition 4, August 04, 2003.

By default, the wavelength is 1550.12 nanometers (nm), which corresponds to 193.40 terahertz (THz).

## SEE ALSO

| *wavelength*

## Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength

To configure the wavelength on 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) and OTN interfaces, include the **wavelength** statement at the **[edit interfaces *interface-name* optics-options]** hierarchy level:

```
[edit interfaces interface-name optics-options]
wavelength nm;
```

To display the currently tuned wavelength and frequency for the interface, use the **show interfaces *interface-name*** operational mode command.

For interface diagnostics, issue the **show interfaces diagnostics optics *interface-name*** operational mode command.

[Table 108 on page 553](#) shows configurable wavelengths and the corresponding frequency for each configurable wavelength.

**Table 108: Wavelength-to-Frequency Conversion Matrix**

| Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) |
|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| 1528.38         | 196.15          | 1542.14         | 194.40          | 1556.15         | 192.65          |
| 1528.77         | 196.10          | 1542.54         | 194.35          | 1556.55         | 192.60          |
| 1529.16         | 196.05          | 1542.94         | 194.30          | 1556.96         | 192.55          |
| 1529.55         | 196.00          | 1543.33         | 194.25          | 1557.36         | 192.50          |

**Table 108: Wavelength-to-Frequency Conversion Matrix (Continued)**

| Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) |
|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| 1529.94         | 195.95          | 1543.73         | 194.20          | 1557.77         | 192.45          |
| 1530.33         | 195.90          | 1544.13         | 194.15          | 1558.17         | 192.40          |
| 1530.72         | 195.85          | 1544.53         | 194.10          | 1558.58         | 192.35          |
| 1531.12         | 195.80          | 1544.92         | 194.05          | 1558.98         | 192.30          |
| 1531.51         | 195.75          | 1545.32         | 194.00          | 1559.39         | 192.25          |
| 1531.90         | 195.70          | 1545.72         | 193.95          | 1559.79         | 192.20          |
| 1532.29         | 195.65          | 1546.12         | 193.90          | 1560.20         | 192.15          |
| 1532.68         | 195.60          | 1546.52         | 193.85          | 1560.61         | 192.10          |
| 1533.07         | 195.55          | 1546.92         | 193.80          | 1561.01         | 192.05          |
| 1533.47         | 195.50          | 1547.32         | 193.75          | 1561.42         | 192.00          |
| 1533.86         | 195.45          | 1547.72         | 193.70          | 1561.83         | 191.95          |
| 1534.25         | 195.40          | 1548.11         | 193.65          | 1562.23         | 191.90          |
| 1534.64         | 195.35          | 1548.51         | 193.60          | 1562.64         | 191.85          |
| 1535.04         | 195.30          | 1548.91         | 193.55          | 1563.05         | 191.80          |

**Table 108: Wavelength-to-Frequency Conversion Matrix (Continued)**

| Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) |
|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| 1535.43         | 195.25          | 1549.32         | 193.50          | 1563.45         | 191.75          |
| 1535.82         | 195.20          | 1549.72         | 193.45          | 1563.86         | 191.70          |
| 1536.22         | 195.15          | 1550.12         | 193.40          | 1564.27         | 191.65          |
| 1536.61         | 195.10          | 1550.52         | 193.35          | 1564.68         | 191.60          |
| 1537.00         | 195.05          | 1550.92         | 193.30          | 1565.09         | 191.55          |
| 1537.40         | 195.00          | 1551.32         | 193.25          | 1565.50         | 191.50          |
| 1537.79         | 194.95          | 1551.72         | 193.20          | 1565.90         | 191.45          |
| 1538.19         | 194.90          | 1552.12         | 193.15          | 1566.31         | 191.40          |
| 1538.58         | 194.85          | 1552.52         | 193.10          | 1566.72         | 191.35          |
| 1538.98         | 194.80          | 1552.93         | 193.05          | 1567.13         | 191.30          |
| 1539.37         | 194.75          | 1553.33         | 193.00          | 1567.54         | 191.25          |
| 1539.77         | 194.70          | 1553.73         | 192.95          | 1567.95         | 191.20          |
| 1540.16         | 194.65          | 1554.13         | 192.90          | 1568.36         | 191.15          |
| 1540.56         | 194.60          | 1554.54         | 192.85          | 1568.77         | 191.10          |

**Table 108: Wavelength-to-Frequency Conversion Matrix (Continued)**

| Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) | Wavelength (nm) | Frequency (THz) |
|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| 1540.95         | 194.55          | 1554.94         | 192.80          |                 |                 |
| 1541.35         | 194.50          | 1555.34         | 192.75          |                 |                 |
| 1541.75         | 194.45          | 1555.75         | 192.70          |                 |                 |

**SEE ALSO**

| [wavelength](#)

**RELATED DOCUMENTATION**

| [Understanding Optical Transport Network \(OTN\) | 458](#)

| [Supported OTN and Optics Options | 491](#)

**Configuring OTN****IN THIS SECTION**

- [Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC | 557](#)
- [Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)
- [Configuring OTN Interfaces on P2-100GE-OTN PIC | 563](#)
- [Configuring OTN Interfaces on MIC3-100G-DWDM MIC | 569](#)
- [Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 575](#)
- [Configuring OTN Interface Options on PTX10K-LC1104 | 581](#)

Use this topic for information about how to configure optical transport network (OTN) interfaces on specific line cards.

## Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC

Starting from Junos OS Release 14.2, a 24-port 10-Gigabit Ethernet OTN PIC—P1-PTX-24-10G-W-SFPP—is supported on the FPC-PTX-P1-A and FPC2-PTX-P1A FPCs in PTX5000 routers, and the FPC-SFF-PTX-P1-A and FPC-SFF-PTX-T FPCs in PTX3000 routers. To configure an OTN interface on the P1-PTX-24-10G-W-SFPP PIC, you must configure interface-specific options and the OTN-related options for the interface.

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
```

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 an 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
```

## SEE ALSO

[Understanding the P1-PTX-24-10G-W-SFPP PIC | 461](#)

[optics-options | 1048](#)

[otn-options | 1051](#)

[signal-degrade](#) | [752](#)

[preemptive-fast-reroute](#) | [747](#)

## Configuring OTN Interfaces on P1-PTX-2-100G-WDM

PTX Series routers support optical transport network (OTN) interfaces, including the 100-Gigabit DWDM OTN PIC (P1-PTX-2-100G-WDM). See *100-Gigabit Ethernet OTN Options Configuration Overview*.

To configure the 100-Gigabit DWDM OTN PIC:

1. Configure the interface wavelength.

```
[edit interfaces interface-name optics-options]
user@host# set wavelength nm
```

See *wavelength*.

**NOTE:** See [100-Gigabit DWDM OTN PIC Integrated Transceiver Optical Interface Specifications](#) for a list of wavelengths supported by the P1-PTX-2-100G-WDM PIC.

2. Enable the laser.

```
[edit interfaces interface-name otn-options]
user@host# set laser-enable
```

3. (Optional) Set the tca.

```
[edit interfaces interface-name otn-options]
user@host# set tca tca-identifier (enable-tca | no-enable-tca) (threshold number |
threshold-24hrs number)
```

See *tca*.

4. (Optional) Set the trace identifiers.

```
[edit interfaces interface-name otn-options]
user@host# set tti tti-identifier tti-identifier-name
```

See *tti*.

5. (Optional) Specify defect triggers.

```
[edit interfaces interface-name otn-options]
user@host# set trigger trigger-identifier
```

See *trigger*.

6. (Optional) Enable VLAN tagging. See *Enabling VLAN Tagging*.
7. (Optional) Set the media MTU. See *Configuring the Media MTU*.
8. (Optional) Set the unit VLAN ID, family `inet`, and IP address.

```
[edit interfaces interface-name]
user@host# set vlan-id number
user@host# set family inet
user@host# set address address
```

9. (Optional) Enable pre-FEC BER signal-degrade monitoring and backward fast reroute to monitor the pre-FEC BER status of the link and to insert the local pre-FEC status into transmitted OTN frames.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set signal-degrade-monitor-enable
user@host# set backward-frr-enable
```

See *signal-degrade-monitor-enable* and *backward-frr-enable*.

10. (Optional) Configure the bit error rate (BER) thresholds for signal degradation used for monitoring the pre-forward error correction (pre-FEC) status of the OTN link.
  - a. Set the BER signal-degrade threshold.

```
[edit interfaces interface-name otn-options signal-degrade]
user@host# set ber-threshold-signal-degrade value
```

- b. Set the BER threshold to clear signal-degrade alarms.

```
[edit interfaces interface-name otn-options signal-degrade]
user@host# set ber-threshold-clear value
```

- c. Set the time interval for signal-degrade collection. After the BER threshold for signal-degrade is crossed for ten consecutive intervals, an alarm is raised. If the BER threshold for signal-degrade

clear is crossed for ten consecutive intervals, the alarm is cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms \* 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

```
[edit interfaces interface-name otn-options signal-degrade]
user@host# set interval value
```

**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.

See *ber-threshold-signal-degrade*, *ber-threshold-clear*, and *interval*.

**NOTE:** See *Understanding Pre-FEC BER Monitoring and BER Thresholds* for more information about pre-FEC BER monitoring and determining BER threshold settings.

## SEE ALSO

*optics-options*

*otn-options*

*signal-degrade*

*preemptive-fast-reroute*

## Configuring OTN Interfaces on P2-100GE-OTN PIC

To configure an OTN interface on the P2-100GE-OTN PIC you must configure interface-specific options and OTN-related options for the interface.

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 VLAN tagging 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
```

To configure the OTN-related options on the interface:

1. Go to the [edit interface *interface-name* *otn-options*] hierarchy level:

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

2. 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
```

3. Set an 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
```

4. 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)
```

5. Enable the threshold crossing alarms for the OTN interface along with the trigger for the defect.

- In Junos OS Release 14.1R2 only:

```
[edit interfaces interface-name otn-options trigger]
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)
```

- In Junos OS Release 14.2 and later:

```
[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)
```

6. 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
```

7. Configure the forward error correction (FEC) mode as Generic Forward Error Correction (GFEC) or none for the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set fec (gfec | none)
```

8. Enable line loopback or local host loopback for the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set line-loopback
user@host# set local-loopback
```

9. Enable an ODU locked maintenance signal on the OTN interface to send the signal pattern 01010101.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-lck
```

10. Enable an ODU open connection indication signal on the OTN interface to send to send the signal pattern 01100110.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-oci
```

11. 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
```

12. 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
```

13. Configure the OTN payload pseudorandom binary sequence (PRBS) on the OTN interface.

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

14. Configure OTN mode as OTU4 for the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set rate otu4
```

15. 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
```

16. 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
```

- ODU backward FRR—Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-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
```

- Monitoring of signal degradation of ODU BER in the received OTN frames.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-signal-degrade-monitor-enable
```

**17.** Configure the following options for ODU BER signal degradation on the OTN interface:

- Configure the threshold for signal degradation for ODU BER when an alarm needs to be raised.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-signal-degrade value
```

- Configure the threshold for ODU BER after signal degradation when the alarm needs to be cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-clear value
```

- When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the ODU 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 ODU BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set interval value
```

## SEE ALSO

[optics-options | 1048](#)

[otn-options | 1051](#)

[signal-degrade | 752](#)

[preemptive-fast-reroute | 747](#)

## Configuring OTN Interfaces on MIC3-100G-DWDM MIC

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. To configure an OTN interface on the MIC3-100G-DWDM MIC, you must configure interface-specific options and OTN-related options for the interface.

To configure the interface-specific options:

1. Configure VLAN tagging at the **[edit interface interface-name]** hierarchy level, where *interface-name* is in the **et-fpc/pic/port** format.

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

2. Configure the maximum transmission unit (MTU) size in bytes for the interface.

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

3. Configure a VLAN ID for the interface.

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

#### 4. Configure the family for the interface.

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

#### 5. Configure an IP address for the interface.

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

To configure the optics-specific options on the interface:

1. 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
```

2. Specify the wavelength of the optics in nanometers. For a list of wavelengths supported, see *wavelength*.

```
[edit interfaces interface-name optics-options]
user@host# set wavelength nm
```

To configure the OTN-specific options on the interface:

1. At the [edit interfaces *interface-name* otn-options] 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
```

2. Set an 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)
```

- By default, triggers are ignored. Specify defect triggers and the set the trigger hold time for the trigger. Possible values for the trigger hold time are as follows: down—Delay before marking interface down when defect occurs (1..65534 milliseconds) and up—Delay before marking interface up when defect is absent (1..65534 milliseconds).

**NOTE:** The hold time value only impacts the alarm reporting time and does not mark an interface down when the defect occurs. To mark the interface up or down, you must also configure the physical interface hold time at the `[edit interfaces interface-name]` hierarchy level.

```
[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)
```

- 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)
```

- 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
```

6. Configure the forward error correction (FEC) mode for the OTN interface. Possible values are: Generic Forward Error Correction (GFEC), or High Gain Forward Error Correction (HGFECC) or Soft Decision Forward Error Correction (SDFEC). The default forward error correction mode is SDFEC.

```
[edit interfaces interface-name otn-options]
user@host# set fec (gfec | hgfec | sdfec)
```

7. Enable line loopback or local host loopback for the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set line-loopback
user@host# set local-loopback
```

8. Enable an ODU locked maintenance signal on the OTN interface to send the signal pattern 01010101.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-lck
```

9. Enable an ODU open connection indication signal on the OTN interface to send to send the signal pattern 01100110.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-oci
```

10. 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
```

11. 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 out-ttim-action-enable
```

12. Configure the OTN payload pseudorandom binary sequence (PRBS) on the OTN interface.

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

13. Configure the line rate or speed of the OTN signal to OTU4 (100Gbps) for the OTN interface.

**NOTE:** If you specify a value other than OTU4, the value is ignored. To verify the line rate, use the **show interfaces *interface-name* extensive** command.

```
[edit interfaces interface-name otn-options]
user@host# set rate otu4
```

14. 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
```

15. 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
```

- ODU backward FRR—Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-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
```

- Monitoring of signal degradation of ODU BER in the received OTN frames.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-signal-degrade-monitor-enable
```

#### 16. Configure the following options for ODU BER signal degradation on the OTN interface:

- Configure the threshold for signal degradation for ODU BER when an alarm needs to be raised.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-signal-degrade value
```

- Configure the threshold for ODU BER after signal degradation when the alarm needs to be cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-clear value
```

- When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the ODU 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 ODU BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set interval value
```

## SEE ALSO

[Understanding the MIC3-100G-DWDM MIC | 474](#)

[optics-options | 1048](#)

[otn-options | 1051](#)

[signal-degrade | 752](#)

[preemptive-fast-reroute | 747](#)

## Configuring OTN Interfaces on PTX-5-100G-WDM PIC

Starting from Junos OS Release 15.1F6, the the 5-port 100-Gigabit DWDM OTN PIC—PTX-5-100G-WDM—is supported on the PTX3000 and the PTX5000 routers. To configure an OTN interface on the PTX-5-100G-WDM PIC, you must configure interface-specific options, optics-specific options and OTN-related options for the interface.

To configure the interface-specific options:

1. Configure VLAN tagging at the [edit interface *interface-name*] hierarchy level, where *interface-name* is in the **et-fpc/pic/port** format.

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

2. Configure the maximum transmission unit (MTU) size in bytes for the interface. Possible values: 256 through 16,000.

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

3. Set the unit VLAN ID, family and the IP address of the interface. Possible values for the VLAN ID: 1 through 4094. Specify the family as **inet**.

```
[edit interfaces interface-name unit 0]
user@host# set vlan-id number
user@host# set family family-name
user@host# set address address
```

To configure the optics-specific options on the interface:

1. 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
```

2. Specify the wavelength of the optics in nanometers. For a list of wavelengths supported, see *wavelength*.

```
[edit interfaces interface-name optics-options]
user@host# set wavelength nm
```

To configure the OTN-specific options on the interface:

1. At the [edit interfaces ***interface-name*** otn-options] hierarchy level, 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
```

2. 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)
```

3. Specify defect triggers and the set the trigger hold time for the trigger. By default, triggers are ignored. Possible values for the trigger hold time are as follows: down and up.
  - down—Delay before marking interface down when defect occurs (1 through 65534 milliseconds)
  - up—Delay before marking interface up when defect is absent (1 through 65534 milliseconds).

**NOTE:** The hold time value only impacts the alarm reporting time and does not mark an interface down when the defect occurs. To mark the interface up or down, you must also configure the physical interface hold time at the `[edit interfaces interface-name]` hierarchy level.

```
[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)
```

4. Enable the threshold-crossing alarms (TCAs) for the OTN interface along with the trigger for the defect. Threshold-crossing alarms (TCAs) are activated when a certain configurable threshold—near-end measurement threshold or far-end measurement threshold—is crossed and remains so until the end of the 15-minute interval for parameters such as OTU and ODU.

```
[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)
```

5. 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
```

6. Configure the forward error correction (FEC) mode for the OTN interface. Possible values are: generic forward error correction (GFEC), or high-gain forward error correction (HG-FEC) or soft-decision forward error correction (SD-FEC). The default forward error correction mode is SD-FEC.

```
[edit interfaces interface-name otn-options]
user@host# set fec (gfec |hgfec | sdfec)
```

7. Enable line loopback or local host loopback for the OTN interface. Loopback testing enables you to verify the connectivity of a circuit. In line loopback, instead of transmitting the signal toward the far-end device, the signal is sent back to the originating router. In local loopback, the signal is transmitted to the channel service unit (CSU) and then to the far-end device.

```
[edit interfaces interface-name otn-options]
user@host# set line-loopback
user@host# set local-loopback
```

8. Enable an ODU locked maintenance signal on the OTN interface to send the signal pattern 01010101.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-lck
```

9. Enable an ODU open connection indication signal on the OTN interface to send the signal pattern 01100110.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-oci
```

10. 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
```

11. 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 out-ttim-action-enable
```

12. Configure the OTN payload pseudorandom binary sequence (PRBS) on the OTN interface.

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

13. Configure the line rate or speed of the OTN signal to otu4 (100 Gbps) for the OTN interface.

**NOTE:** If you specify a value other than otu4, the value is ignored. To verify the line rate, use the **show interfaces *interface-name* extensive** command.

```
[edit interfaces interface-name otn-options]
user@host# set rate otu4
```

14. 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
```

15. 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
```

- ODU backward FRR—Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-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
```

- Monitoring of signal degradation of ODU BER in the received OTN frames.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-signal-degrade-monitor-enable
```

**16.** Configure the following options for ODU BER signal degradation on the OTN interface:

- Configure the threshold for signal degradation for ODU BER when an alarm needs to be raised.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-signal-degrade value
```

- Configure the threshold for ODU BER after signal degradation when the alarm needs to be cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-clear value
```

- When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the ODU 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 ODU BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set interval value
```

## SEE ALSO

*Before You Begin Installing or Upgrading the Firmware*

*Installing Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*

[Understanding the PTX-5-100G-WDM PIC | 478](#)

*Upgrading Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*

[optics-options | 1048](#)

[otn-options | 1051](#)

[signal-degrade | 752](#)

[preemptive-fast-reroute | 747](#)

## Configuring OTN Interface Options on PTX10K-LC1104

The PTX10K-LC1104 line card provides up to 1.2 Tbps packet forwarding for cloud providers, service providers, and enterprises that need coherent dense wavelength-division multiplexing (DWDM) with MACsec security features. The PTX10K-LC1104 line card is supported on Junos OS Release 18.3R1 and later.

Each PTX10K-LC1104 has 6 physical interfaces (ot-x/x/x) that connect to one of three built-in flexible rate optical transponders. Each transponder connects four 100-Gigabit Ethernet logical interfaces (et-x/x/x) to one of three forwarding ASICs.

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. 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. For a list of wavelengths supported, see *wavelength*.

```
[edit interfaces interface-name optics-options]
user@host# set wavelength nm
```

To configure the OTN-specific options on the interface:

1. At the `[edit interfaces interface-name otn-options]` 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
```

2. Set an 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)
```

3. By default, triggers are ignored. Specify defect triggers and the set the trigger hold time for the trigger. Possible values for the trigger hold time are as follows: down—Delay before marking interface down when defect occurs (1..65534 milliseconds) and up—Delay before marking interface up when defect is absent (1..65534 milliseconds).

**NOTE:** The hold time value only impacts the alarm reporting time and does not mark an interface down when the defect occurs. To mark the interface up or down, you must also configure the physical interface hold time at the `[edit interfaces interface-name]` hierarchy level.

```
[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)
```

4. 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)
```

5. 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
```

6. Configure the forward error correction (FEC) mode for the OTN interface. Possible values are: Generic Forward Error Correction (GFEC), or High Gain Forward Error Correction (HGFECC) or Soft Decision Forward Error Correction (SDFEC). The default forward error correction mode is SDFEC.

```
[edit interfaces interface-name otn-options]
user@host# set fec (gfec | hgfec | sdfec)
```

7. Enable line loopback or local host loopback for the OTN interface.

```
[edit interfaces interface-name otn-options]
user@host# set line-loopback
user@host# set local-loopback
```

8. Enable an ODU locked maintenance signal on the OTN interface to send the signal pattern 01010101.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-lck
```

9. Enable an ODU open connection indication signal on the OTN interface to send to send the signal pattern 01100110.

```
[edit interfaces interface-name otn-options]
user@host# set insert-odu-oci
```

10. 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
```

11. 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 out-ttim-action-enable
```

12. Configure the OTN payload pseudorandom binary sequence (PRBS) on the OTN interface.

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

13. Configure the line rate or speed of the OTN signal to OTU4 (100Gbps) for the OTN interface.

**NOTE:** If you specify a value other than OTU4, the value is ignored. To verify the line rate, use the **show interfaces *interface-name* extensive** command.

```
[edit interfaces interface-name otn-options]
user@host# set rate otu4
```

14. 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
```

15. 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
```

- ODU backward FRR—Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-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
```

- Monitoring of signal degradation of ODU BER in the received OTN frames.

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
user@host# set odu-signal-degrade-monitor-enable
```

## 16. Configure the following options for ODU BER signal degradation on the OTN interface:

- Configure the threshold for signal degradation for ODU BER when an alarm needs to be raised.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-signal-degrade value
```

- Configure the threshold for ODU BER after signal degradation when the alarm needs to be cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set ber-threshold-clear value
```

- When you configure the interval along with the **ber-threshold-signal-degrade *value*** statement, the ODU 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 ODU BER must stay below the clear threshold for the configured interval after which the alarm is cleared.

```
[edit interfaces interface-name otn-options odu-signal-degrade]
user@host# set interval value
```

## SEE ALSO

[optics-options | 1048](#)

[otn-options | 1051](#)

[signal-degrade](#) | [752](#)

[preemptive-fast-reroute](#) | [747](#)

[Understanding the PTX10K-LC1104 Line Card](#) | [482](#)

#### Release History Table

| Release | Description                                                                                                                                        |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| 15.1F6  | Starting from Junos OS Release 15.1F6, the the 5-port 100-Gigabit DWDM OTN PIC—PTX-5-100G-WDM—is supported on the PTX3000 and the PTX5000 routers. |

#### RELATED DOCUMENTATION

[Understanding Optical Transport Network \(OTN\)](#) | [458](#)

[Supported OTN and Optics Options](#) | [491](#)

## ODU Path Delay Measurement for Performance Monitoring

#### IN THIS SECTION

- [Understanding ODU Path Delay Measurement](#) | [587](#)
- [Enabling ODU Path Delay Measurement](#) | [589](#)
- [Disabling ODU Path Delay Measurement](#) | [591](#)

Use this topic to understand about ODU path delay measurement and performance monitoring.

### Understanding ODU Path Delay Measurement

#### IN THIS SECTION

- [Guidelines for Configuring Delay Measurement](#) | [588](#)

Performance monitoring is an important requirement in any network, including the optical transport networks (OTN). 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 back to you; the time taken for the round-trip indicates the latency.

The optical channel data unit (ODU) path delay measurement offers in-service delay measurement. Delay (or latency) is measured by transmitting a known pattern (delay measurement pattern) in a selected bit of the delay measurement (**DM**) field and measuring the number of frames that are missed when the delay measurement pattern is received at the transmitting end. For instance, if the transmitted delay measurement bit is **1 1 1 1 1 1 1 1 0 0** and the received delay measurement bit is **1 1 1 0 0 0 0 0 0 0**, the delay measurement starts at frame 2 and ends at frame 8. This can be detected by the change in value between the transmitted bit and the received bit.

|           |    |   |   |   |   |   |   |   |   |   |
|-----------|----|---|---|---|---|---|---|---|---|---|
| Frame#    | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| Tx DM bit | 1  | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 0 |
| Rx DM bit | 1  | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

The result of the delay measurement is 6 frames (8 - 2).

### Guidelines for Configuring Delay Measurement

When you configure in-service delay measurement, we recommend that you follow certain guidelines to ensure that you obtain accurate 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.

**NOTE:** 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.

## SEE ALSO

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit DWDM OTN MIC with CFP2-ACO](#)

[100-Gigabit DWDM OTN PIC with CFP2-ACO \(PTX Series\)](#)

[Configuring OTN Interfaces on MIC3-100G-DWDM MIC | 569](#)

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 575](#)

[remote-loop-enable | 750](#)

[Understanding the MIC3-100G-DWDM MIC | 474](#)

[Understanding the PTX-5-100G-WDM PIC | 478](#)

## 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, first enable remote loopback of the delay measurement pattern on the remote interface and then start measurement of the delay.

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
...
```

## SEE ALSO

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit DWDM OTN MIC with CFP2-ACO](#)

[100-Gigabit DWDM OTN PIC with CFP2-ACO \(PTX Series\)](#)

[Configuring OTN Interfaces on MIC3-100G-DWDM MIC | 569](#)

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 575](#)

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 575](#)

[remote-loop-enable | 750](#)

[Understanding the MIC3-100G-DWDM MIC | 474](#)

[Understanding the PTX-5-100G-WDM PIC | 478](#)

## 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
...
```

**SEE ALSO**

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

---

[100-Gigabit DWDM OTN MIC with CFP2-ACO](#)

---

[100-Gigabit DWDM OTN PIC with CFP2-ACO \(PTX Series\)](#)

---

[Configuring OTN Interfaces on MIC3-100G-DWDM MIC | 569](#)

---

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 575](#)

---

[remote-loop-enable | 750](#)

---

[Understanding the MIC3-100G-DWDM MIC | 474](#)

---

[Understanding the PTX-5-100G-WDM PIC | 478](#)

**RELATED DOCUMENTATION**

[Understanding Optical Transport Network \(OTN\) | 458](#)

---

[Supported OTN and Optics Options | 491](#)

# 3

PART

## Troubleshooting Information

---

[Monitoring and Troubleshooting Ethernet Interfaces](#) | 595

---

# Monitoring and Troubleshooting Ethernet Interfaces

## IN THIS CHAPTER

- [Passive Monitoring on Ethernet Interfaces Overview | 595](#)
- [Enabling Passive Monitoring on Ethernet Interfaces | 597](#)
- [Link Degrade Monitoring Overview | 600](#)
- [Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 604](#)
- [Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 621](#)
- [Performing Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 643](#)
- [Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 663](#)
- [Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters | 671](#)
- [Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State | 679](#)
- [Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test | 681](#)

## 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.

#### Release History Table

| 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 595](#). 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 | 595](#)

[Configuring Multiservice Physical Interface Properties](#)

[Junos OS Services Interfaces Library for Routing Devices](#)

[Ethernet Interfaces User Guide for Routing Devices](#)

## Link Degradation Monitoring Overview

### IN THIS SECTION

- [Supported Platforms | 601](#)

Link degradation 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 degradation 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 degradation 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 109 on page 602](#) and [Table 110 on page 603](#) lists the platform series and line cards that support link degrade monitoring.

Table 109: MX Line Cards that Support Link Degradation Monitoring

| Platform Series | MPC Line Cards Supported                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | DPC Line Cards Supported                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MX              | <ul style="list-style-type: none"> <li>• MPC4E-3D-2CGE-8XGE</li> <li>• MPC4E-3D-32XGE-SFPP</li> <li>• MPC-3D-16XGE-SFP</li> <li>• MPC3 with MIC3-3D-1X100GE-CFP</li> <li>• MPC3 with MIC3-3D-2X40GE-QSFPP</li> <li>• MPC3 with MIC-3D-2XGE-XFP</li> <li>• MPC3 with 2x10GE XFP MIC</li> <li>• MPC3 with 2x10GE XFP MIC</li> <li>• MPC5 with following variants:               <ul style="list-style-type: none"> <li>• 2CGE + 4XGE</li> <li>• 24XGE + 6XLGE</li> </ul> </li> <li>• MPC6 with the following variants:               <ul style="list-style-type: none"> <li>• 2X100GE CFP2</li> <li>• 24X10GE SFPP</li> <li>• 24X10GE SFPP OTN</li> <li>• 4x100GE CXP</li> </ul> </li> <li>• MPC7E-MRATE</li> </ul> | <ul style="list-style-type: none"> <li>• DPCE-R-Q-4XGE-XFP</li> <li>• DPCE-R-4XGE-XFP</li> <li>• DPCE-X-4XGE-XFP</li> <li>• DPCE-X-Q-4XGE-XFP</li> <li>• DPCE-R-2XGE-XFP</li> <li>• DPCE-R-4XGE-XFP</li> <li>• DPCE-X-4XGE-XFP</li> </ul> <p data-bbox="1040 842 1321 909">On 10-Gigabit Ethernet interfaces:</p> <ul style="list-style-type: none"> <li>• DPCE-R-Q-20GE-2XGE</li> <li>• DPCE-R-20GE-2XGE</li> <li>• DPCE-X-20GE-2XGE</li> </ul> |

Table 109: MX Line Cards that Support Link Degradation Monitoring (*Continued*)

| Platform Series | MPC Line Cards Supported                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | DPC Line Cards Supported |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
|                 | <ul style="list-style-type: none"> <li>MPC7E-10G (non-MACsec mode)</li> <li>MX2K-MPC8E with MIC-MRATE</li> <li>MX2K-MPC9E with MIC-MRATE</li> </ul> <p><b>NOTE:</b> Link degradation monitoring is not supported on MACsec-enabled MPC7E-10G and MIC-MACSEC-MRATE.</p> <ul style="list-style-type: none"> <li>MPC10E-10C-MRATE (10G, 25G, 40G, 100G and 400G interfaces)</li> <li>MPC10E-15C-MRATE (10G, 25G, 40G, 100G and 400G interfaces)</li> <li>MX2K-MPC11E (10G, 40G, 100G, and 400G interfaces)</li> </ul> |                          |

Table 110: PTX Line Cards that Support Link Degradation Monitoring

| Platform Series             | Line Cards/PICs Supported                                                                                        |
|-----------------------------|------------------------------------------------------------------------------------------------------------------|
| PTX3000 and PTX5000 (Junos) | <ul style="list-style-type: none"> <li>P1-PTX-2-100GE-CFP</li> <li>P2-100GE-CFP2</li> <li>PF-1CGE-CFP</li> </ul> |
| PTX10003 (Junos Evolved)    | 10G, 25G, 40G, 50G, 100G, 200G, and 400G interfaces                                                              |

**Table 110: PTX Line Cards that Support Link Degrade Monitoring (Continued)**

| Platform Series          | Line Cards/PICs Supported                                                                                                 |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------|
| PTX10004 (Junos Evolved) | 10G, 25G, 40G, 50G, 100G, 200G, and 400G interfaces                                                                       |
| PTX10008 (Junos Evolved) | <ul style="list-style-type: none"> <li>• JNP10K-LC1201</li> <li>• PTX10K-LC1202-36MR (10G and 100G interfaces)</li> </ul> |

**RELATED DOCUMENTATION**

[Physical Interface Damping Overview](#)

[Fast Reroute Overview](#)

[link-degrade-monitor | 955](#)

[thresholds | 1179](#)

[recovery | 1101](#)

[request interface link-degrade-recover | 1243](#)

**Monitoring Fast Ethernet and Gigabit Ethernet Interfaces****IN THIS SECTION**

- [Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 605](#)
- [Monitor Fast Ethernet and Gigabit Ethernet Interfaces | 606](#)
- [Fiber-Optic Ethernet Interface Specifications | 619](#)

## Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

### IN THIS SECTION

- Purpose | 605
- Action | 605
- Meaning | 606

### Purpose

To monitor Fast Ethernet and Gigabit Ethernet interfaces and begin the process of isolating interface problems when they occur.

### Action

[Table 111 on page 605](#) provides links and commands for monitoring Fast Ethernet and Gigabit Ethernet interfaces.

**Table 111: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces**

| Tasks                                                                                                     | Command or Action                                                          |
|-----------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| <i>Monitor Fast Ethernet and Gigabit Ethernet Interfaces</i>                                              |                                                                            |
| 1. <i>Display the Status of Fast Ethernet Interfaces</i>                                                  | <b>show interfaces terse (fe*   ge*)</b>                                   |
| 1. <i>Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface</i>                    | <b>show interfaces (fe- fpc/ pic/ port   ge- fpc/ pic/ port)</b>           |
| 1. <i>Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface</i> | <b>show interfaces (fe- fpc/ pic/ port   ge- fpc/ pic/ port) extensive</b> |
| 1. <i>Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface</i>                            | <b>monitor interface (fe- fpc/ pic/ port   ge- fpc/ pic/ port)</b>         |

**Table 111: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces (Continued)**

| Tasks                                                   | Command or Action |
|---------------------------------------------------------|-------------------|
| 1. <i>Fiber-Optic Ethernet Interface Specifications</i> |                   |

**Meaning**

You can use the above described commands to monitor and to display the configurations for Fast Ethernet and Gigabit Ethernet interfaces.

**SEE ALSO**

*Display the Status of Gigabit Ethernet Interfaces*

*Display the Status of Fast Ethernet Interfaces*

**Monitor Fast Ethernet and Gigabit Ethernet Interfaces****IN THIS SECTION**

- [Display the Status of Fast Ethernet Interfaces | 607](#)
- [Display the Status of Gigabit Ethernet Interfaces | 609](#)
- [Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface | 610](#)
- [Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface | 612](#)
- [Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 617](#)

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 | 607
- Action | 607
- Meaning | 608

### *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 up inet 10.116.115.217/29
fe-3/0/2 up down
fe-3/0/2.0 up down
fe-3/0/3 up up
fe-3/0/3.0 up up inet 192.168.223.65/30
fe-4/1/0 down up
fe-4/1/0.0 up down inet 10.150.59.133/30
fe-4/1/1 up up
fe-4/1/1.0 up up inet 10.150.59.129/30
fe-4/1/2 up down
fe-4/1/2.0 up down
```

*Meaning*

The sample output lists only the 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 112: Status of Fast Ethernet Interfaces**

| Physical Interface                | Logical Interface                   | Status Description                                                                                                                                                                                                  |
|-----------------------------------|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fe-2/1/0<br>Admin Up<br>Link Up   | fe-2/1/0.0<br>Admin Up<br>Link Up   | This interface has both the physical and logical links up and running.                                                                                                                                              |
| fe-3/0/2<br>Admin Up<br>Link Down | fe-3/0/2.0<br>Admin Up<br>Link Down | This interface has the physical link down, the link layer down, or both down ( <b>Link Down</b> ). The logical link is also down as a result.                                                                       |
| fe-4/1/0<br>Admin Down<br>Link Up | fe-4/1/0.0<br>Admin Up<br>Link Down | This interface is administratively disabled and the physical link is healthy ( <b>Link Up</b> ), but the logical interface is not established. The logical interface is down because the physical link is disabled. |
| fe-4/1/2<br>Admin Up<br>Link Down | fe-4/1/2.0<br>Admin Up<br>Link Down | This interface has both the physical and logical links down.                                                                                                                                                        |

**SEE ALSO**

| [Monitoring Fast Ethernet and Gigabit Ethernet Interfaces](#) | 604

## Display the Status of Gigabit Ethernet Interfaces

### IN THIS SECTION

- Purpose | 609
- Action | 609
- Meaning | 609

### *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 interfaces terse ge*
Interface Admin Link Proto Local Remote
ge-2/2/0 down down
ge-2/2/0.0 up down inet 65.113.23.105/30
ge-2/3/0 up up
ge-2/3/0.0 up up inet 65.115.56.57/30
ge-3/1/0 up up
ge-3/1/0.0 up up inet 65.115.56.193/30
ge-3/2/0 up down
```

### *Meaning*

This sample output lists only the Gigabit Ethernet interfaces. It shows the status of both the physical and logical interfaces. See [Table 3](#) for a description of what the output means.

Table 113: Status of Gigabit Ethernet Interfaces

| Physical Interface                  | Logical Interface                   | Status Description                                                                                                                    |
|-------------------------------------|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| ge-2/2/0<br>Admin Down<br>Link Down | ge-2/2/0.0<br>Admin Up<br>Link Down | This interface is administratively disabled ( <b>Admin Down</b> ). Both the physical and logical links are down ( <b>Link Down</b> ). |
| ge-2/3/0<br>Admin Up<br>Link Up     | ge-2/3/0.0<br>Admin Up<br>Link Up   | This interface has both the physical and logical links up and running.                                                                |
| ge-3/2/0<br>Admin Up<br>Link Down   | ge-3/2/0.0<br>Admin Up<br>Link Down | This interface has both the physical link and the logical interface down.                                                             |

## SEE ALSO

[Monitoring Fast Ethernet and Gigabit Ethernet Interfaces](#) | 604

## Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface

## IN THIS SECTION

- [Purpose](#) | 611
- [Action](#) | 611
- [Meaning](#) | 612

### *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*, 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 | 612
- Action | 613
- Meaning | 614

### ***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
 MAC statistics:

```

|                   | Receive   | Transmit  |
|-------------------|-----------|-----------|
| Total octets      | 439703575 | 177452093 |
| Total packets     | 1866532   | 1642916   |
| Unicast packets   | 972137    | 1602563   |
| Broadcast packets | 30        | 2980      |
| Multicast packets | 894365    | 37373     |
| CRC/Align errors  | 0         | 0         |

```

FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Filter statistics:
Input packet count 1866532
Input packet rejects 0
Input DA rejects 503674
Input SA rejects 0
Output packet count 1642916
Output packet pad count 0
Output packet error count 0
CAM destination filters: 5, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete, Link partner status: OK
Link partner: Full-duplex, Flow control: None
PFE configuration:
Destination slot: 1, Stream number: 15
CoS transmit queue bandwidth:
 Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
CoS weighted round-robin:
 Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
Description: Test
Flags: SNMP-Traps, Encapsulation: ENET2
Protocol inet, MTU: 1500, Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.115.107.192/29, Local: 10.115.107.193
 Broadcast: 10.115.107.199

```

### **Meaning**

The sample output shows where the errors might be occurring and includes autonegotiation information. See [Table 4](#) for a description of errors to look for.

**Table 114: Errors to Look For**

| Error                    | Meaning                                                                                                                                                                   |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Policed discards</b>  | Discarded frames that were not recognized or were not of interest.                                                                                                        |
| <b>L2 channel errors</b> | 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.       |
| <b>MTU</b>               | 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. |
| <b>Input DA rejects</b>  | 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.                        |
| <b>Input SA rejects</b>  | 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* for an explanation of Fast Ethernet and Gigabit Ethernet alarms.

[Table 5](#) lists and describes some MAC statistics errors to look for.

**Table 115: MAC Statistics Errors**

| Error                     | Meaning                                                                                                                                                                                                                                                                                                           |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>CRC/Align errors</b>   | 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 ( <b>FCS Error</b> ) or a bad FCS with a non-integral number of octets ( <b>Alignment Error</b> ). |
| <b>MAC control frames</b> | The number of MAC control frames.                                                                                                                                                                                                                                                                                 |

Table 115: MAC Statistics Errors (*Continued*)

| Error                   | Meaning                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MAC pause frames</b> | The number of MAC control frames with <b>pause</b> operational code.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Jabber frames</b>    | <p>The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error.</p> <p>Note that this definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.</p> |
| <b>Fragment frames</b>  | <p>The total number of packets received that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error.</p> <p>Note that it is entirely normal for fragment frames to increment because both runts (which are normal occurrences due to collisions) and noise hits are counted.</p>                                                                                                                      |

Autonegotiation is the process that connected Ethernet interfaces use to communicate the information necessary to interoperate. [Table 6](#) explains the autonegotiation information of the **show interface *interface-name* extensive** command output.

Table 116: Autonegotiation Information

| Autonegotiation Field Information             | Explanation                                                                                                                                                                        |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Negotiation status: Incomplete</b>         | The <b>Negotiation status</b> field shows <b>Incomplete</b> when the Ethernet interface has the speed or link mode configured.                                                     |
| <b>Negotiation status: No autonegotiation</b> | The <b>Negotiation status</b> field shows <b>No autonegotiation</b> when the remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. |

Table 116: Autonegotiation Information (*Continued*)

| Autonegotiation Field Information                                     | Explanation                                                                                                                                                                                                                                            |
|-----------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Negotiation status: Complete</b><br><b>Link partner status: OK</b> | The <b>Negotiation status</b> field shows <b>Complete</b> and the <b>Link partner</b> field shows <b>OK</b> when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process completes successfully. |
| <b>Link partner: Half-duplex</b>                                      | The <b>Link partner</b> field can be <b>Full-duplex</b> or <b>Half-duplex</b> depending on the capability of the attached Ethernet device.                                                                                                             |
| <b>Flow control: Symmetric/asymmetric</b>                             | The <b>Flow control</b> 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 | 617
- Action | 617
- Meaning | 618

#### *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:

```

user@host> monitor interface fe-2/1/0
Interface: fe-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100mbps
Traffic statistics:
Input bytes: 282556864218 (14208 bps) [40815]
Output bytes: 42320313078 (384 bps) [890]
Input packets: 739373897 (11 pps) [145]
Output packets: 124798688 (1 pps) [14]
Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
Input framing errors: 0 [0]
Policed discards: 6625892 [6]
L3 incompletes: 75 [0]
L2 channel errors: 0 [0]
L2 mismatch timeouts: 0 [0]
Carrier transitions: 1 [0]
Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
Unicast packets 464751787 [154]
Packet error count 0 [0]

```

## Meaning

Use the information from this command to help narrow down possible causes of an interface problem.

**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* 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* for the optical specifications.

## Fiber-Optic Ethernet Interface Specifications

Table 117 on page 619 shows the specifications for fiber-optic interfaces for Juniper Networks routers.

**Table 117: Fiber-Optic Ethernet Interface Specifications**

| Fiber-Optic Ethernet Interface | Length                                      | Wavelength      | Average Launch Power | Receiver Saturation | Receiver Sensitivity                     |
|--------------------------------|---------------------------------------------|-----------------|----------------------|---------------------|------------------------------------------|
| <b>Gigabit Ethernet</b>        |                                             |                 |                      |                     |                                          |
| <b>Duplex SC connector</b>     |                                             |                 |                      |                     |                                          |
| <b>LH optical interface</b>    | 49.5-mile 70-km reach on 8.2-micrometer SMF | 1480 to 1580 nm | -3 to +2 dBm         | -3 dBm              | -23 dBm (BER 10 <sup>-12</sup> ) for SMF |

Table 117: Fiber-Optic Ethernet Interface Specifications (Continued)

| Fiber-Optic Ethernet Interface                   | Length                                                                                                           | Wavelength      | Average Launch Power | Receiver Saturation | Receiver Sensitivity |
|--------------------------------------------------|------------------------------------------------------------------------------------------------------------------|-----------------|----------------------|---------------------|----------------------|
| <b>LX optical interface</b>                      | 6.2-mile 10-km reach on 9/125-micrometer SMF<br><br>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              |
| <b>SX optical interface</b>                      | 656-ft 200-m reach on 62.5/125-micrometer MMF<br><br>1640-ft 500-m reach on 50/125-micrometer MMF                | 830 to 860 nm   | -9.5 to -4 dBm       | -3 dBm              | -17 dBm              |
| <b>Fast Ethernet 8-Port</b>                      |                                                                                                                  |                 |                      |                     |                      |
| <b>FX optical interface with MT-RJ connector</b> | 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 | 621
- Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface | 623
- Create a Loopback | 624
- Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 627
- Configure a Static Address Resolution Protocol Table Entry | 632
- Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 637
- Ping the Fast Ethernet or Gigabit Ethernet Interface | 637
- Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 639
- Diagnose a Suspected Circuit Problem | 642

## Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

### IN THIS SECTION

- Purpose | 621
- Action | 621

#### Purpose

To use loopback testing to isolate Fast Ethernet and Gigabit Ethernet interface problems.

#### Action

[Table 118 on page 622](#) provides links and commands for using loopback testing for Fast Ethernet and Gigabit Ethernet interfaces.

Table 118: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

| Tasks                                                                                           | Command or Action                                                                                                                                                                                                                                   |
|-------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface</i> |                                                                                                                                                                                                                                                     |
| 1. <i>Create a Loopback</i>                                                                     |                                                                                                                                                                                                                                                     |
| 1. a. <i>Create a Physical Loopback for a Fiber-Optic Interface</i>                             | Connect the transmit port to the receive port.                                                                                                                                                                                                      |
| 1. a. <i>Create a Loopback Plug for an RJ-45 Ethernet Interface</i>                             | Cross pin 1 (TX+) and pin 3 (RX+) together, and pin 2 (TX-) and pin 6 (RX-) together.                                                                                                                                                               |
| 1. a. <i>Configure a Local Loopback</i>                                                         | <code>[edit interfaces <i>interface-name</i> (fastether-options   gigheter-options)] set loopback show commit</code>                                                                                                                                |
| 1. <i>Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up</i>                     | <code>show interfaces (fe-<i>fpc/pic/port</i>   ge-<i>fpc/pic/port</i>)</code>                                                                                                                                                                      |
| 1. <i>Configure a Static Address Resolution Protocol Table Entry</i>                            | <code>show interfaces ge-<i>fpc/pic/port</i> [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>] set arp <i>ip-address</i> mac <i>mac-address</i> show commit run show arp no-resolve</code> |
| 1. <i>Clear Fast Ethernet or Gigabit Ethernet Interface Statistics</i>                          | <code>clear interfaces statistics fe-<i>fpc/pic/port</i> / ge-<i>fpc/pic/port</i></code>                                                                                                                                                            |
| 1. <i>Ping the Fast Ethernet or Gigabit Ethernet Interface</i>                                  | <code>ping <i>remote-IP-address</i> bypass-routing interface (fe-<i>fpc/pic/port</i>   ge-<i>fpc/pic/port</i>) count 100 rapid</code>                                                                                                               |

**Table 118: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces**  
(Continued)

| Tasks                                                                            | Command or Action                                                                                                                |
|----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| 1. <i>Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics</i> | <b>show interfaces (fe-<i>fpc/pic/port</i>   ge-<i>fpc/pic/port</i>) extensive</b>                                               |
| <i>Diagnose a Suspected Circuit Problem</i>                                      | Perform Steps 2 through 8 from <i>Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface</i> . |

## Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface

### IN THIS SECTION

- Problem | 623
- Solution | 623

### 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*
- *Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up*
- *Configure a Static Address Resolution Protocol Table Entry*
- *Clear Fast Ethernet or Gigabit Ethernet Interface Statistics*

- *Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics*

## Create a Loopback

### IN THIS SECTION

- [Create a Physical Loopback for a Fiber-Optic Interface | 624](#)
- [Create a Loopback Plug for an RJ-45 Ethernet Interface | 625](#)
- [Configure a Local Loopback | 626](#)

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 | 621](#)

## 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 8: RJ-45 Ethernet Loopback Plug

### 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](#) | 621

## 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:

```
[edit interfaces interface-name (fastether-options | gigether-options)]
user@host# set 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 | gigheter-options)]** hierarchy level, places the interface into local loopback mode.

**NOTE:** Remember to delete the loopback statement after completing the test.

## SEE ALSO

[Configure a Local Loopback | 0](#)

## RELATED DOCUMENTATION

[Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 627](#)

[Configure a Static Address Resolution Protocol Table Entry | 632](#)

[Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 637](#)

[Ping the Fast Ethernet or Gigabit Ethernet Interface | 637](#)

[Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 639](#)

## Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up

### IN THIS SECTION

- [Purpose | 627](#)
- [Action | 628](#)
- [Meaning | 631](#)

### 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
Physical interface: ge-4/0/6, Enabled, Physical link is Up Interface index: 144, SNMP ifIndex:
516, Generation: 147
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags : Present Running Loop-Detected
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 0
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
Last flapped : 2015-01-20 23:40:04 PST (00:02:12 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets : 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets : 0
Dropped traffic statistics due to STP State:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets : 0
Input errors:
```

```

Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO
errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0
Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control

Active alarms : None
Active defects : None
MAC statistics:
Total octets Receive Transmit
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0

```

```

 CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Complete
 Link partner:
 Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote
fault: OK
 Local resolution:
 Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 4
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit % bps % usec
0 best-effort 95 950000000 95 0 low
none
3 network-control 5 50000000 5 0 low
none
Interface transmit statistics: Disabled

```

## Meaning

The sample output shows that the link is up and there are no alarms in this loopback configuration. When an internal loopback is configured, the physical loopback should come up without an alarm.

## Sample Output

When you see that the physical link is down, there may be a problem with the port. The following output is an example of the `show interfaces fe-fpc/pic/port` command when the physical link is down:

```

user@router> show interfaces fe-1/3/0
Physical interface: fe-1/3/0, Enabled, Physical link is Down
 Interface index: 44, SNMP ifIndex: 35
 Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
 Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
 Device flags : Present Running Down
 Interface flags: Hardware-Down SNMP-Traps
 Link flags : None
 Current address: 00:90:69:8d:2c:db, Hardware address: 00:90:69:8d:2c:db
 Input rate : 0 bps (0 pps), Output rate: 0 bps (0 pps)
Active alarms : LINK

```

**Active defects : LINK**

## MAC statistics:

Input octets: 0, Input packets: 0, Output octets: 0, Output packets: 0

## Filter statistics:

Filtered packets: 0, Padded packets: 0, Output packet errors: 0

## Autonegotiation information:

Negotiation status: Incomplete, Link partner status: Down

Reason: Link partner autonegotiation failure

Link partner: Half-duplex, Flow control: None

**Meaning**

The sample output shows that the physical link is down and there are active alarms and defects.

[Table 119 on page 631](#) presents problem situations and actions for a physical link that is down.

**Table 119: 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
 Link flags : None
 CoS queues : 8 supported, 4 maximum usable queues
 Schedulers : 0
 Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
 Last flapped : 2015-01-20 23:40:04 PST (00:13:49 ago)
 Statistics last cleared: 2015-01-20 23:46:15 PST (00:07:38 ago)
 Traffic statistics:
 Input bytes : 125500 0 bps
 Output bytes : 125482 0 bps
 Input packets: 1281 0 pps
 Output packets: 1281 0 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
```

```

Output packets: 0
Dropped traffic statistics due to STP State:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO
errors: 0,
 Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped
packets
 0 best-effort 1260
1260 0
 1 expedited-fo 0
0 0
 2 assured-forw 0
0 0
 3 network-cont 0
0 0
Queue number: Mapped forwarding classes
 0 best-effort
 1 expedited-forwarding
 2 assured-forwarding
 3 network-control
Active alarms : None
Active defects : None
MAC statistics:
Total octets Receive Transmit
Total packets 1281 1281
Unicast packets 1280 1280
Broadcast packets 1 1
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0

```

```

Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Filter statistics:
 Input packet count 1281
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 1281
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Complete
 Link partner:
 Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote
fault: OK
 Local resolution:
 Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
 Destination slot: 4
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer
Priority Limit % bps % usec
 0 best-effort 95 950000000 95 0
low none
 3 network-control 5 50000000 5 0
low none
Interface transmit statistics: Disabled

Logical interface ge-4/0/6.0 (Index 72) (SNMP ifIndex 573) (Generation 137)
Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 125500
 Output bytes : 123480
 Input packets: 1281
 Output packets: 1260
Local statistics:
 Input bytes : 60

```

```

Output bytes : 2002
Input packets: 1
Output packets: 21
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp
nhrp ospf pgm
pim rip router-discovery rsvp sap vrrp
Flow Statistics :
Flow Input statistics :
 Self packets : 0
 ICMP packets : 40
 VPN packets : 0
 Multicast packets : 0
 Bytes permitted by policy : 107520
 Connections established : 20
Flow Output statistics:
 Multicast packets : 0
 Bytes permitted by policy : 107520
Flow error statistics (Packets dropped due to):
 Address spoofing: 0
 Authentication failed: 0
 Incoming NAT errors: 0
 Invalid zone received packet: 0
 Multiple user authentications: 0
 Multiple incoming NAT: 0
 No parent for a gate: 0
 No one interested in self packets: 0
 No minor session: 0
 No more sessions: 0
 No NAT gate: 0
 No route present: 11
 No SA for incoming SPI: 0
 No tunnel found: 0
 No session for a gate: 0
 No zone or NULL zone binding 0
 Policy denied: 0
 Security association not active: 0
 TCP sequence number out of window: 0

```

```

 Syn-attack protection: 0
 User authentication errors: 0
 Protocol inet, MTU: 1500, Generation: 158, Route table: 0
 Flags: Sendbroadcast-pkt-to-re
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.108.120.0/30, Local: 10.108.120.1, Broadcast:
10.108.120.3,
 Generation: 158
 Protocol multiservice, MTU: Unlimited, Generation: 159, Route table: 0
 Policer: Input: __default_arp_policer__

```

2. In configuration mode, go to the following hierarchy level:

```

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

```

3. Configure the static ARP entry:

```

user@host# set arp ip-address mac mac-address

```

4. Commit the configuration:

```

user@host# commit

```

5. Verify that the static ARP entry is installed:

```

[edit interfaces ge-4/0/6 unit 0 family inet address 10.108.120.1/30]
user@host# run show arp no-resolve

```

| MAC Address       | Address       | Interface  | Flags     |
|-------------------|---------------|------------|-----------|
| 00:1f:12:fe:c5:2e | 10.108.120.2  | ge-4/0/6.0 | permanent |
| 52:54:00:7d:33:4c | 10.204.128.35 | fxp0.0     | none      |
| 52:54:00:65:11:50 | 10.204.128.36 | fxp0.0     | none      |
| 52:54:00:da:30:82 | 10.204.128.37 | fxp0.0     | none      |
| 52:54:00:3a:cf:4b | 10.204.128.38 | fxp0.0     | none      |

```
52:54:00:de:88:5f 10.204.128.45 fxp0.0 none
52:54:00:48:03:b7 10.204.128.46 fxp0.0 none
```

### Meaning

The sample output is for Step 1 through Step 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:

```
user@host> clear interfaces statistics (fe-fpc/pic/port | ge-fpc/pic/
port)
```

### Sample Output

```
user@host> clear interfaces statistics ge-4/0/6
user@host>
```

### Meaning

This command clears the interface statistics counters for the Gigabit Ethernet interface only.

## Ping the Fast Ethernet or Gigabit Ethernet Interface

### IN THIS SECTION

- Purpose | 638
- Action | 638
- Meaning | 638

## Purpose

Use the ping command to verify the loopback connection.

## Action

To send ping packets from the Ethernet interface, use the following Junos OS CLI operational mode command:

```
user@host> ping remote-IP-address bypass-routing interface (fe-fpc/pic/port | ge-
fpc/pic/port) count 100 rapid
```

## Sample Output

```
user@router> ping 10.108.120.2 bypass-routing interface ge-7/2/1 count 100 rapid
PING 10.108.120.2 (10.108.120.2): 56 data bytes
36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
 4 5 00 0054 e871 0 0000 01 01 cc5c 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
 4 5 00 0054 e874 0 0000 01 01 cc59 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
 4 5 00 0054 e878 0 0000 01 01 cc55 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
 4 5 00 0054 e87c 0 0000 01 01 cc51 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
 4 5 00 0054 e880 0 0000 01 01 cc4d 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
Vr HL TOS Len ID Flg off TTL Pro cks Src Dst
 4 5 00 0054 e884 0 0000 01 01 cc49 10.108.120.1 10.108.120.2
.36 bytes from 10.108.120.1: Time to live exceeded
```

## Meaning

The sample output shows that the time to live (TTL) expired, indicating that the link is receiving 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 TTL 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 packet is either lost, or the TTL 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 | 639
- Action | 639
- Meaning | 642

### Purpose

Persistent interface error statistics indicate that you need to open a case with the Juniper Networks Technical Assistance Center (JTAC).

### Action

To check the local interface for error statistics, use the following Junos OS CLI operational mode command:

```
user@host> show interfaces (fe-fpc/pic/port | ge-fpc/pic/port)
extensive
```

## Sample Output

```

user@router> show interfaces ge-4/0/6 extensive
Physical interface: ge-4/0/6, Enabled, Physical link is Up Interface index: 144, SNMP ifIndex:
516, Generation: 147
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags : Present Running Loop-Detected
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 0
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
Last flapped : 2015-01-20 23:40:04 PST (00:02:12 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets : 0 0 pps
Output packets : 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets : 0
Dropped traffic statistics due to STP State:
Input bytes : 0
Output bytes : 0
Input packets : 0
Output packets : 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO
errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets

```

```

0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0
Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control

Active alarms : None
Active defects : None

MAC statistics:
Total octets Receive Transmit
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0

Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote
fault: OK
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:

```

```

Destination slot: 4
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority
Limit % bps % usec
0 best-effort 95 950000000 95 0 low
none
3 network-control 5 50000000 5 0 low
none
Interface transmit statistics: Disabled

```

## Meaning

Check for any error statistics. There should not be any input or output errors. If there are any persistent input or output errors, open a case with the Juniper Networks Technical Assistance Center (JTAC) at [support@juniper.net](mailto: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 | 642
- Action | 642

### 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*. 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 | 643](#)
- [Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | 645](#)
- [Create a Loopback | 646](#)
- [Verify That the Gigabit Ethernet Interface Is Up | 648](#)
- [Configure a Static Address Resolution Protocol Table Entry | 654](#)
- [Clear Gigabit Ethernet Interface Statistics | 655](#)
- [Ping the Gigabit Ethernet Interface | 656](#)
- [Check for Gigabit Ethernet Interface Error Statistics | 658](#)
- [Diagnose a Suspected Circuit Problem | 662](#)

## Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces

### IN THIS SECTION

- [Purpose | 643](#)
- [Action | 644](#)

### Purpose

To use loopback testing to isolate 10, 40, and 100 Gigabit Ethernet interface problems.

## Action

Table 120 on page 644 provides links and commands for using loopback testing for 10, 40, and 100 Gigabit Ethernet interfaces.

**Table 120: Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces**

| Tasks                                                                          | Command or Action                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface</i> |                                                                                                                                                                                                                                                                     |
| 1. Create a Loopback                                                           |                                                                                                                                                                                                                                                                     |
| 1. Create a Physical Loopback for a Fiber-Optic Interface                      | Connect the transmit port to the receive port.                                                                                                                                                                                                                      |
| 1. Configure a Local Loopback                                                  | [edit interfaces <i>interface-name</i> (gigether-options)] set loopback show commit                                                                                                                                                                                 |
| 1. Verify That the Gigabit Ethernet Interface Is Up                            | show interfaces (xe- <i>fpc/pic/port</i>   et- <i>fpc/pic/port</i> )                                                                                                                                                                                                |
| 1. Configure a Static Address Resolution Protocol Table Entry                  | show interfaces (xe- <i>fpc/pic/port</i>   et- <i>fpc/pic/port</i> )[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> ]set arp <i>ip-address</i> mac <i>mac-address</i> show commit run show arp no-resolve |
| 1. Clear Gigabit Ethernet Interface Statistics                                 | clear interfaces statistics xe- <i>fpc/pic/port</i> / et- <i>fpc/pic/port</i>                                                                                                                                                                                       |
| 1. Ping the Gigabit Ethernet Interface                                         | ping <i>remote-IP-address</i> bypass-routing interface (xe- <i>fpc/pic/port</i>   et- <i>fpc/pic/port</i> ) count 100                                                                                                                                               |

**Table 120: Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces**  
(Continued)

| Tasks                                                           | Command or Action                                                                                                                |
|-----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| 1. <i>Check for Gigabit Ethernet Interface Error Statistics</i> | <b>show interfaces (xe-<i>fpc/pic/port</i>   et-<i>fpc/pic/port</i>) extensive</b>                                               |
| <i>Diagnose a Suspected Circuit Problem</i>                     | Perform Steps 2 through 8 from <i>Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface</i> . |

## Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface

### IN THIS SECTION

- Problem | 645
- Solution | 645

### 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*
- *Verify That the Gigabit Ethernet Interface Is Up*
- *Configure a Static Address Resolution Protocol Table Entry*
- *Clear Gigabit Ethernet Interface Statistics*
- *Check for Gigabit Ethernet Interface Error Statistics*

## Create a Loopback

### IN THIS SECTION

- [Create a Physical Loopback for a Fiber-Optic Interface | 646](#)
- [Configure a Local Loopback | 647](#)

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

---

[Create a Physical Loopback for a Fiber-Optic Interface | 0](#)

---

[Verify That the Gigabit Ethernet Interface Is Up | 648](#)

---

[Configure a Static Address Resolution Protocol Table Entry | 654](#)

## 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 gigheter-options
```

2. Configure the local loopback:

```
[edit interfaces interface-name gigheter-options]
user@router# set loopback
```

3. Verify the configuration:

```
user@router# show
```

For example:

```
[edit interfaces xe-2/0/0 gigheter-options]
user@router# show
loopback;
```

4. Commit the change:

```
user@router# commit
```

For example:

```
[edit interfaces xe-2/0/0 gigheter-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 | 648](#)

[Configure a Static Address Resolution Protocol Table Entry | 654](#)

## RELATED DOCUMENTATION

[Verify That the Gigabit Ethernet Interface Is Up | 648](#)

[Configure a Static Address Resolution Protocol Table Entry | 654](#)

[Clear Gigabit Ethernet Interface Statistics | 655](#)

[Ping the Gigabit Ethernet Interface | 656](#)

[Check for Gigabit Ethernet Interface Error Statistics | 658](#)

## Verify That the Gigabit Ethernet Interface Is Up

### IN THIS SECTION

● [Purpose | 648](#)

● [Action | 649](#)

● [Meaning | 653](#)

### 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
 Hold-times : Up 4000 ms, Down 0 ms
 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 : 0
 Output bytes : 0
 Input packets : 0
 Output packets : 0
 Dropped traffic statistics due to STP State:
 Input bytes : 0
```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
 L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
 Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
 0 6400 6400 0
 1 0 0 0
 2 0 0 0
 3 0 0 0
Queue number: Mapped forwarding classes
 0 best-effort
 1 expedited-forwarding
 2 assured-forwarding
 3 network-control
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 rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 6400
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

```

Packet Forwarding Engine configuration:

Destination slot: 0 (0x00)

CoS information:

Direction : Output

| CoS transmit queue | Bandwidth | Buffer     | Priority |      |     |
|--------------------|-----------|------------|----------|------|-----|
| Limit              | %         | bps        | %        | usec |     |
| 0 best-effort      | 95        | 9500000000 | 95       | 0    | low |
| none               |           |            |          |      |     |
| 3 network-control  | 5         | 500000000  | 5        | 0    | low |
| none               |           |            |          |      |     |

Preclassifier statistics:

| Traffic Class   | Received Packets | Transmitted Packets | Dropped |
|-----------------|------------------|---------------------|---------|
| real-time       | 0                | 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
 Policers: 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-fpc/pic/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 : 0
 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 0
Interface transmit statistics: Disabled

```

### Meaning

The sample output shows that the physical link is down and there are active alarms and defects.

[Table 121 on page 653](#) presents problem situations and actions for a physical link that is down.

**Table 121: 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. |

**Table 121: Problems and Solutions for a Physical Link That Is Down (Continued)**

| Problem                                      | Action                                                                                                                                                                                                                                                                        |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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:

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

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
```

| MAC Address       | Address      | Interface  | Flags     |
|-------------------|--------------|------------|-----------|
| 02:01:00:00:00:05 | 10.0.0.5     | em1.0      | none      |
| 00:00:5e:00:01:01 | 10.85.175.1  | fxp0.0     | none      |
| d8:b1:22:0a:6e:00 | 10.85.175.2  | fxp0.0     | none      |
| d0:07:ca:57:d7:a0 | 10.85.175.3  | fxp0.0     | none      |
| 00:a0:a5:c2:06:e2 | 10.85.175.4  | fxp0.0     | none      |
| d8:18:d3:b3:6d:ea | 10.108.120.2 | xe-2/0/0.0 | permanent |
| 02:01:00:00:00:05 | 128.0.0.5    | em1.0      | none      |
| 02:01:00:00:00:05 | 128.0.0.6    | em1.0      | none      |
| 02:00:00:00:00:10 | 128.0.0.16   | em0.0      | none      |
| 02:00:00:00:00:12 | 128.0.0.18   | em0.0      | none      |
| 02:00:00:00:00:17 | 128.0.0.23   | em0.0      | none      |
| 02:00:00:00:00:1a | 128.0.0.26   | em0.0      | none      |
| Total 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:

```
user@router> clear interfaces statistics (xe-fpc/pic/port | et-
fpc/pic/port)
```

### 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 | 656
- Action | 656
- Meaning | 657

### Purpose

Use the ping command to verify the loopback connection.

### Action

To send ping packets from the Ethernet interface, use the following Junos OS CLI operational mode command:

```
user@router> ping remote-IP-address bypass-routing interface (xe-fpc/pic/port | et-
fpc/pic/port) count 100 rapid
```

## 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 Src 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 Src 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 Src 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 Src 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 Src 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 Src 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 TTL 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 packet is either lost, or the TTL 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 | [658](#)
- Action | [658](#)
- Meaning | [661](#)

### Purpose

Persistent interface error statistics indicate that you need to open a case with the Juniper Networks Technical Assistance Center (JTAC).

### Action

To check the local interface for error statistics, use the following Junos OS CLI operational mode command:

```
user@router> (xe-fpc/pic/port | et-fpc/pic/port) extensive
```

### 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
Hold-times : Up 4000 ms, Down 0 ms
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 : 0
 Output packets: 0
Dropped traffic statistics due to STP State:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
 L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
 Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
 0 6400 6400 0
 1 0 0 0
 2 0 0 0
 3 0 0 0
Queue number: Mapped forwarding classes
 0 best-effort
 1 expedited-forwarding
 2 assured-forwarding
 3 network-control
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 rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 6400
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
 Destination slot: 0 (0x00)
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 0 best-effort 95 9500000000 95 0 low
none
 3 network-control 5 500000000 5 0 low
none
Preclassifier statistics:
 Traffic Class Received Packets Transmitted Packets Dropped
Packets
 real-time 0

```

```

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
 Policers: 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](mailto: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 | 662

● Action | 662

### 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*. 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 | 663](#)
- [Configuring BERT Testing | 666](#)
- [Starting and Stopping a BERT Test | 670](#)

### 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, ADSL, 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 122 on page 664 shows the loopback modes supported on the various interface types.

**Table 122: Loopback Modes by Interface Type**

| Interface                                            | Loopback Modes   | Usage Guidelines                                                                                                                                                                                                                                                                                                                      |
|------------------------------------------------------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet | Local            | <a href="#">Configuring Ethernet Loopback Capability</a>                                                                                                                                                                                                                                                                              |
| Circuit Emulation E1                                 | Local and remote | <a href="#">Configuring E1 Loopback Capability</a>                                                                                                                                                                                                                                                                                    |
| Circuit Emulation T1                                 | Local and remote | <a href="#">Configuring T1 Loopback Capability</a>                                                                                                                                                                                                                                                                                    |
| E1 and E3                                            | Local and remote | <a href="#">Configuring E1 Loopback Capability</a> and <a href="#">Configuring E3 Loopback Capability</a>                                                                                                                                                                                                                             |
| NxDS0                                                | Payload          | <a href="#">Configuring NxDS0 IQ and IQE Interfaces</a> , <a href="#">Configuring T1 and NxDS0 Interfaces</a> , <a href="#">Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode)</a> , <a href="#">Configuring Fractional E1 IQ and IQE Interfaces</a> , and <a href="#">Configuring Channelized T3 IQ Interfaces</a> |

Table 122: Loopback Modes by Interface Type (*Continued*)

| Interface              | Loopback Modes                           | Usage Guidelines                                                                                                                                                      |
|------------------------|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Serial (V.35 and X.21) | Local and remote                         | <a href="#">Configuring Serial Loopback Capability</a>                                                                                                                |
| Serial (EIA-530)       | DCE local, DCE remote, local, and remote | <a href="#">Configuring Serial Loopback Capability</a>                                                                                                                |
| SONET/SDH              | Local and remote                         | <a href="#">Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External</a>                                                               |
| T1 and T3              | Local, payload, and remote               | <a href="#">Configuring T1 Loopback Capability and Configuring T3 Loopback Capability</a><br><br>See also <a href="#">Configuring the T1 Remote Loopback Response</a> |

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)
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-2e7 Pattern is 2^7 - 1
pseudo-2e9-o153 Pattern is 2^9 - 1 (per 0.153 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:** 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 O.152 and O.153 (2047
type) standards)
pseudo-2e15-o151 Pattern is 2^15 -1 (per O.151 standard)
pseudo-2e20-o151 Pattern is 2^20 -1 (per O.151 standard)
pseudo-2e20-o153 Pattern is 2^20 -1 (per O.153 standard)
pseudo-2e23-o151 Pattern is 2^23 -1 (per O.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
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
pseudo-2e9-o153 Pattern is 2^9 - 1 (per O.153 standard)
pseudo-2e11-o152 Pattern is 2^11 - 1 (per O.152 standard)
pseudo-2e15-o151 Pattern is 2^15 - 1 (per O.151 standard)
pseudo-2e20-o151 Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e20-o153 Pattern is 2^20 - 1 (per O.153 standard)
pseudo-2e23-o151 Pattern is 2^23 (per O.151 standard)

```

Table 123 on page 669 shows the BERT capabilities for various interface types.

**Table 123: BERT Capabilities by Interface Type**

| Interface                                              | T1 BERT            | T3 BERT                     | Comments                                                                                                                                                                      |
|--------------------------------------------------------|--------------------|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 12-port T1/E1<br>Circuit Emulation                     | Yes (ports 0-11)   | —                           | <ul style="list-style-type: none"> <li>Limited algorithms</li> </ul>                                                                                                          |
| 4-port<br>Channelized<br>OC3/STM1<br>Circuit Emulation | Yes (port 0-3)     | —                           | <ul style="list-style-type: none"> <li>Limited algorithms</li> </ul>                                                                                                          |
| E1 or T1                                               | Yes (port 0-3)     | Yes (port 0-3)              | <ul style="list-style-type: none"> <li>Single port at a time</li> <li>Limited algorithms</li> </ul>                                                                           |
| E3 or T3                                               | Yes (port 0-3)     | Yes (port 0-3)              | <ul style="list-style-type: none"> <li>Single port at a time</li> </ul>                                                                                                       |
| Channelized<br>OC12                                    | —                  | Yes (channel 0-11)          | <ul style="list-style-type: none"> <li>Single channel at a time</li> <li>Limited algorithms</li> <li>No bit count</li> </ul>                                                  |
| Channelized<br>STM1                                    | Yes (channel 0-62) | —                           | <ul style="list-style-type: none"> <li>Multiple channels</li> <li>Only one algorithm</li> <li>No error insert</li> <li>No bit count</li> </ul>                                |
| Channelized T3<br>and Multichannel<br>T3               | Yes (channel 0-27) | Yes (port 0-3 on channel 0) | <ul style="list-style-type: none"> <li>Multiple ports and channels</li> <li>Limited algorithms for T1</li> <li>No error insert for T1</li> <li>No bit count for T1</li> </ul> |

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 \(Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port\) | 1368](#)

## 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 | 671](#)
- [Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm | 672](#)
- [Fast Ethernet and Gigabit Ethernet Counters | 675](#)

## Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

### IN THIS SECTION

- [Purpose | 671](#)
- [Action | 671](#)

### Purpose

To locate LINK alarm and major counters associated with Fast Ethernet and Gigabit Ethernet interfaces.

### Action

[Table 124 on page 672](#) provides links and commands for locating LINK alarm and major counters for Fast Ethernet and Gigabit Ethernet interfaces.

**Table 124: Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters**

| Tasks                                                                     | Command or Action                                                    |
|---------------------------------------------------------------------------|----------------------------------------------------------------------|
| <i>Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm</i> | <b>show interfaces (fe-fpc/pic/port   ge-fpc/pic/port) extensive</b> |
| <i>Fast Ethernet and Gigabit Ethernet Counters</i>                        |                                                                      |

**SEE ALSO**

| [Ethernet Interfaces User Guide for Routing Devices](#)

**Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm****IN THIS SECTION**

- [Problem | 672](#)
- [Solution | 672](#)

**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**

```
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@host> show interfaces fe-1/3/3 extensive
Physical interface: fe-1/3/3, Enabled, Physical link is Down
 Interface index: 47, SNMP ifIndex: 38
 Description: Test
 Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
 Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link flags : None
 Current address: 00:90:69:8d:2c:de, Hardware address: 00:90:69:8d:2c:de
 Statistics last cleared: 2002-01-11 23:03:09 UTC (1w2d 23:54 ago)
 Traffic statistics:
 Input bytes : 373012658 0 bps
 Output bytes : 153026154 1392 bps
 Input packets : 1362858 0 pps
 Output packets : 1642918 3 pps
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 503660
 L3 incompletes: 1 , L2 channel errors: 0, L2 mismatch timeouts: 0
 FIFO errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Collisions: 0, Drops: 0, Aged packets: 0
 HS link CRC errors: 0, FIFO errors: 0

Active alarms : LINK
Active defects : LINK
MAC statistics:

```

|                    | Receive   | Transmit  |
|--------------------|-----------|-----------|
| Total octets       | 439703575 | 177452093 |
| Total packets      | 1866532   | 1642916   |
| Unicast packets    | 972137    | 1602563   |
| Broadcast packets  | 30        | 2980      |
| Multicast packets  | 894365    | 37373     |
| CRC/Align errors   | 0         | 0         |
| FIFO errors        | 0         | 0         |
| MAC control frames | 0         | 0         |
| MAC pause frames   | 0         | 0         |
| Oversized frames   | 0         |           |
| Jabber frames      | 0         |           |
| Fragment frames    | 0         |           |
| VLAN tagged frames | 0         |           |
| Code violations    | 0         |           |

```

Filter statistics:
 Input packet count 1866532
 Input packet rejects 0
 Input DA rejects 503674
 Input SA rejects 0
 Output packet count 1642916
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 5, CAM source filters: 0
Autonegotiation information:
 Negotiation status: Complete, Link partner status: OK
 Link partner: Full-duplex, Flow control: None
PFE configuration:
 Destination slot: 1, Stream number: 15
 CoS transmit queue bandwidth:
 Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
 CoS weighted round-robin:
 Queue0: 95, Queue1: 0, Queue2: 0, Queue3: 5
Logical interface fe-1/3/3.0 (Index 8) (SNMP ifIndex 69)
 Description: Test
 Flags: SNMP-Traps, Encapsulation: ENET2
 Protocol inet, MTU: 1500, Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.115.107.192/29, Local: 10.115.107.193
 Broadcast: 10.115.107.199

```

## Meaning

The sample output shows where the alarm and other errors might be occurring and any counters that are incrementing. The only alarm associated with 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* 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

[Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters | 671](#)

## Fast Ethernet and Gigabit Ethernet Counters

### IN THIS SECTION

● Problem | 675

● Solution | 675

### Problem

#### Description

Table 125 on page 675 shows the major counters that appear in the output for the **show interfaces fe-*fpc/pic/port* extensive** and the **show interfaces ge-*fpc/pic/port* extensive** commands. These counters generally increment when there is a problem with a 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 125: Major Fast Ethernet and Gigabit Ethernet Counters

| Counter                 | Description                                                                     | Reason for Increment                                                                                                                                |
|-------------------------|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input Errors:</b>    |                                                                                 |                                                                                                                                                     |
| <b>Errors</b>           | The sum of the incoming frame terminates and frame check sequence (FCS) errors. |                                                                                                                                                     |
| <b>Policed discards</b> | 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 125: Major Fast Ethernet and Gigabit Ethernet Counters (Continued)

| Counter                     | Description                                                                                                                            | Reason for Increment                                                                                                                                                                                                         |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Drops</b>                | 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.                                                                            |
| <b>L3 incompletes</b>       | 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. |
| <b>L2 channel errors</b>    | The errors that occur when the software could not find a valid logical interface (such as <b>fe-1/2/3.0</b> ) for an incoming frame.   | This error increments when, for example, a lookup for a virtual LAN (VLAN) fails.                                                                                                                                            |
| <b>L2 mismatch timeouts</b> | The count of malformed or short packets.                                                                                               | The malformed or short packets cause the incoming packet handler to discard the frame and be unreadable.                                                                                                                     |
| <b>FIFO errors</b>          | 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.                                                                                          |
| <b>Output Errors</b>        |                                                                                                                                        |                                                                                                                                                                                                                              |
| <b>Errors</b>               | The sum of outgoing frame terminates and FCS errors.                                                                                   |                                                                                                                                                                                                                              |

Table 125: Major Fast Ethernet and Gigabit Ethernet Counters (Continued)

| Counter                                | Description                                                                                                                                        | Reason for Increment                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Collisions</b>                      | 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.                                                                                                                                                                                                                          |
| <b>Drops</b>                           | 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.                                                                                                                                                                                                                                                   |
| <b>Aged packets</b>                    | 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.                                                                                                                                                                                                                                                        |
| <b>HS link FCS errors, FIFO errors</b> | 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.                                                                                                                                                                                                                                                                         |
| <b>Miscellaneous Counters</b>          |                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Input DA rejects</b>                | 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). |

Table 125: Major Fast Ethernet and Gigabit Ethernet Counters (*Continued*)

| Counter                                            | Description                                                                                                                        | Reason for Increment                                                                                                                                                                                                                                                                  |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Output packet pad count</b>                     | 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. |
| <b>Output packet error count</b>                   | Number of packets with an indicated error that the filter was given to transmit.                                                   | These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.                                                                                                                |
| <b>CAM destination filters, CAM source filters</b> | 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 | 671](#)

[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 | [679](#)
- Diagnosis | [679](#)
- Resolution | [680](#)

### 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 680](#) 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 | 681](#)
- [Guidelines to perform Bidirectional Diagnostics using Remote Loopback | 686](#)
- [Clearing the Interface Statistics | 689](#)

### 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 <b>show interfaces prbs-stats</b> 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 <b>show interfaces prbs-stats</b> displays data for four lanes. |

*(Continued)*

| Interface Type                 | Interface Name | Interface Lane Characteristics                           | Remarks                                                                                                                      |
|--------------------------------|----------------|----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| 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 <b>show interfaces prbs-stats</b> 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-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
```

2. Start a RX (direction 1) by issuing the following command:

```
[edit]
user@host2> test interface et-1/1/4 prbs-test-start pattern-type 31 direction 1 flip 0
```

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 0
```

For example:

To stop the PRBS at RX:

```
user@host2> test interface et-1/1/4 prbs-test-stop direction 1
```

Verify the statistics at TX by executing the following command:

```
show interfaces et-0/1/2 prbs-stats
PRBS Statistics : Disabled
```

To stop the PRBS at TX:

```
user@host1> test interface et-0/1/2 prbs-test-stop direction 0
```

After executing the command, you can check the link status at RX by executing **show interfaces terse et-1/1/4**:

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|-------|-------|--------|
| et-1/1/4  | up    | up   |       |       |        |

Check the link status at TX by executing **show interfaces terse et-0/1/2**:

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|-------|-------|--------|
| et-0/1/2  | up    | up   |       |       |        |

This command only disables the statistics collection and does not clear the statistics collected. To clear the collected statistics, issue the **clear interfaces statistics** command.

```
user@host1> clear interfaces statistics et-0/1/2
```

## SEE ALSO

[prbs-test-start | 1239](#)

[prbs-test-stop | 1242](#)

[show interfaces prbs-stats | 1703](#)

*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 gigeother-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 gigeother-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-MIC1-MACSEC 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 gigheter-options loopback-remote
set interfaces et-1/1/1 gigheter-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 values are INT\_MAX, the delta value 0 is displayed.

Consider a scenario where the error count increments as the PRBS test is ongoing. In this case, the **show interfaces *interface-name* prbs-stats** shows the incremental error (delta value). Also after issuing **clear interfaces statistics *et-1/1/4***, if the error count is updated, the **clear interfaces statistics** displays the incremental error too.

For Example:

```
[edit]
user@host2> show interfaces et-1/1/4 prbs-stats
PRBS Statistics : Enabled
 Lane 0 : State : Fail, Error count : 640
 Lane 1 : State : Fail, Error count : 647
 Lane 2 : State : Fail, Error count : 661
 Lane 3 : State : Fail, Error count : 596

[edit]
user@host2> show interfaces et-1/1/4 prbs-stats
PRBS Statistics : Enabled
 Lane 0 : State : Fail, Error count : 52
 Lane 1 : State : Fail, Error count : 65
 Lane 2 : State : Fail, Error count : 626
 Lane 3 : State : Fail, Error count : 132
```

As shown above, in the second instance the **show interfaces *interface-name* prbs-stats** command displays the "change" in error count. So, the total error count is  $640 + 52 = 692$  for Lane 0.

## RELATED DOCUMENTATION

[prbs-test-start | 1239](#)

---

[prbs-test-stop | 1242](#)

---

[show interfaces prbs-stats | 1703](#)

---

[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test | 681](#)

# 4

PART

## Configuration Statements and Operational Commands

---

[Configuration Statements \(OTN\) | 693](#)

[Configuration Statements | 777](#)

[Operational Commands | 1228](#)

---

# Configuration Statements (OTN)

## IN THIS CHAPTER

- alarm (optics-options) | 694
- backward-frr-enable | 696
- ber-threshold-clear | 698
- ber-threshold-signal-degrade | 701
- bypass | 704
- bytes (otn-options) | 706
- cfp-to-et | 707
- fec | 709
- fec | 712
- fec (gigether) | 713
- fixed-stuff-bytes | 715
- high-polarization | 717
- interval | 718
- is-ma | 721
- laser-enable | 722
- line-loopback | 724
- local-loopback | 725
- monitor-end-point | 727
- no-odu-backward-frr-enable | 728
- no-odu-signal-degrade-monitor-enable | 729
- number-of-frames | 731
- oc192 | 732
- odu-delay-management | 733
- odu-backward-frr-enable | 735
- odu-signal-degrade | 736
- odu-signal-degrade-monitor-enable | 738

- [odu-ttim-action-enable | 739](#)
- [otu-ttim-action-enable | 741](#)
- [otu4 | 742](#)
- [pass-through | 744](#)
- [prbs | 745](#)
- [preemptive-fast-reroute | 747](#)
- [rate | 748](#)
- [remote-loop-enable | 750](#)
- [signal-degrade | 752](#)
- [signal-degrade-monitor-enable | 753](#)
- [start-measurement | 755](#)
- [tca | 757](#)
- [transport-monitoring | 761](#)
- [trigger | 762](#)
- [tti | 767](#)
- [tx-power | 768](#)
- [warning | 770](#)
- [wavelength | 771](#)

## alarm (optics-options)

### IN THIS SECTION

- [Syntax | 695](#)
- [Hierarchy Level | 695](#)
- [Description | 695](#)
- [Options | 695](#)
- [Required Privilege Level | 695](#)
- [Release Information | 695](#)

## Syntax

```
alarm low-light-alarm {
 (link-down | syslog);
}
```

## Hierarchy Level

```
[edit interfaces interface-name optics-options]
[edit dynamic-profiles name interfaces name optics-options],
[edit dynamic-profiles name logical-systems name interfaces name
optics-options],
```

## Description

Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold. Starting in Junos OS Release 15.1, for all QSFP-based interfaces, you need not explicitly configure the **syslog** option. The **syslog** option is enabled by default.

## Options

**low-light-alarm**—Enable the low light alarm.

**link-down**—Drop the 10-Gigabit Ethernet link and marks link as down.

**syslog**—Write the optics information to the system log.

## 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 Release 10.0.

## RELATED DOCUMENTATION

[Configuring Link Down Notification for Optics Options Alarm or Warning | 214](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

## backward-frr-enable

### IN THIS SECTION

- [Syntax | 696](#)
- [Hierarchy Level | 696](#)
- [Description | 696](#)
- [Default | 697](#)
- [Options | 697](#)
- [Required Privilege Level | 697](#)
- [Release Information | 697](#)

### Syntax

```
(backward-frr-enable | no-backward-frr-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Description

Enable or disable backward fast reroute status insertion.

Enable backward fast reroute to insert local pre-forward error correction (FEC) bit error rate (BER) status into transmitted OTN frames, notifying the remote interface. The remote interface can use the information to reroute traffic to a different interface. When you enable backward fast reroute and also enable pre-FEC BER monitoring including the [signal-degrade-monitor-enable](#) statement,

notification of signal degradation and rerouting of traffic occurs in less time than that required through a Layer 3 protocol.

**NOTE:** When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the **signal-degrade-monitor-enable** and **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. Include the `ber-threshold-signal-degrade value`, `ber-threshold-clear value`, and `interval value` statements at the `[edit interfaces interface-name otn-options signal-degrade]` hierarchy level to configure the BER thresholds and time interval. See "[Understanding Pre-FEC BER Monitoring and BER Thresholds](#)" on page 543 for more information about pre-FEC BER monitoring and determining BER threshold settings.

## Default

By default, backward fast reroute insertion is disabled.

## Options

**backward-frr-enable**—Enable backward fast reroute status insertion.

**no-backward-frr-enable**—Do not enable backward fast reroute status insertion.

## 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 Release 13.2.

## RELATED DOCUMENTATION

[Understanding Pre-FEC BER Monitoring and BER Thresholds](#) | 543

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 459

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 561

## ber-threshold-clear

### IN THIS SECTION

- [Syntax | 698](#)
- [Hierarchy Level | 698](#)
- [Description | 698](#)
- [Options | 700](#)
- [Required Privilege Level | 700](#)
- [Release Information | 700](#)

### Syntax

```
ber-threshold-clear value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

### Description

Specify bit error rate (BER) threshold to clear the interface alarm for signal degradation.

You can configure the BER clear threshold to customize the BER that will clear an interface alarm when signal degrade monitoring is enabled.

**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.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options signal-degrade]` hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-

FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See ["Understanding Pre-FEC BER Monitoring and BER Thresholds" on page 543](#) for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

**NOTE:** You can configure ODU BER thresholds only at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level on the P2-100GE-OTN PIC.

[Table 126 on page 699](#) shows the default values for pre-FEC BER and ODU BER signal degrade threshold values for different PICs. If the BER signal degrade threshold is not configured, the default value is used.

**Table 126: Default Clear Threshold Values**

| PIC or MPC           | Default Pre-FEC BER Clear Threshold Value | Default ODU BER Clear Threshold Value |
|----------------------|-------------------------------------------|---------------------------------------|
| P1-PTX-2-100G-WDM    | 3.0E-3                                    | Not supported                         |
| P2-100GE-OTN         | 3.0E-3                                    | 1.0E-9                                |
| P1-PTX-24-10G-W-SFPP | 3.0E-3                                    | Not supported                         |
| MIC6-100G-CFP2       | 1.0E-6                                    | 1.0E-9                                |
| MPC5E                | 1.0E-6                                    | 1.0E-9                                |

To configure the threshold that raises the signal degrade alarm, include the `ber-threshold-signal-degrade` statement at the same hierarchy level. To configure the time interval during which the BER must stay above or below the configured thresholds to raise or clear the alarm, include the `interval` statement at the same hierarchy level.

**NOTE:** For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below

the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms \* 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

## Options

- **Values:** *value*—BER threshold for clearing the signal degradation in scientific notation. Both the mantissa and exponent are configurable. Enter the value in the format  $xE-n$ , where  $x$  is the mantissa and  $n$  is the exponent. For example, 4.5E-3.

- **Range:** The mantissa must be a decimal number. There is no limit on the number of digits before or after the decimal point.

The exponent must be an integer from 0 through 9.

- **Default:** See [Table 126 on page 699](#) for the default values.

**BEST PRACTICE:** Always set the **ber-threshold-clear** *value* lower than the **ber-threshold-signal-degrade** *value*. For the FEC limits, see the table describing the signal degrade and clear thresholds after configuration in "[Understanding Pre-FEC BER Monitoring and BER Thresholds](#)" on page 543.

**NOTE:** In Junos OS Release 13.2R1, only the exponent is valid input for the BER threshold value, and the mantissa is not configurable. The BER threshold value is  $1.0E-n$  where  $n > 0$ , and the valid range of  $n$  is from 1 through 10.

## 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 Release 13.2.

## RELATED DOCUMENTATION

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 543](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## ber-threshold-signal-degrade

### IN THIS SECTION

- [Syntax | 701](#)
- [Hierarchy Level | 701](#)
- [Description | 701](#)
- [Options | 703](#)
- [Required Privilege Level | 704](#)
- [Release Information | 704](#)

### Syntax

```
ber-threshold-signal-degrade value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

### Description

Specify the bit error rate (BER) threshold to raise an interface alarm for signal degradation.

You can configure the BER signal degrade threshold to customize the BER that will raise an interface alarm when signal degrade monitoring is enabled.

**NOTE:** Configuring a high BER threshold for signal degradation and a long interval might cause the internal bit error counter register to get saturated. For example, for the P1-PTX-2-100G-WDM PIC, the internal bit error counter gets saturated when the error count reaches  $2E+29$ . Therefore, the value of `ber-threshold-signal-degrade * line rate / interval` must be less than  $2E+29$  to avoid saturation. Assuming a fixed PIC line rate of  $1.27E+11$  bits per second and an interval of 1000 ms, the **ber-threshold-signal-degrade** value must be less than  $4.22E-3$ . If the value of the `ber-threshold-signal-degrade * line rate / interval` exceeds the saturation limit, the configuration is ignored by the router, and the default values are used instead. A system log message is logged for this error.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options signal-degrade]` hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See "[Understanding Pre-FEC BER Monitoring and BER Thresholds](#)" on page 543 for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

**NOTE:** You can configure ODU BER thresholds only at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level on the P2-100GE-OTN PIC.

[Table 127 on page 702](#) shows the default values for pre-FEC BER and ODU BER signal degrade threshold values for different PICs. If the BER signal degrade threshold is not configured, the default value is used.

**Table 127: Default Signal Degrade Threshold Values**

| PIC or MPC        | Default Pre-FEC BER Signal Degrade Threshold Value | Default ODU BER Signal Degrade Threshold Value |
|-------------------|----------------------------------------------------|------------------------------------------------|
| P1-PTX-2-100G-WDM | $7.5E-3$                                           | Not supported                                  |
| P2-100GE-OTN      | $7.5E-3$                                           | $1.0E-6$                                       |

Table 127: Default Signal Degrade Threshold Values (Continued)

| PIC or MPC           | Default Pre-FEC BER Signal Degrade Threshold Value | Default ODU BER Signal Degrade Threshold Value |
|----------------------|----------------------------------------------------|------------------------------------------------|
| P1-PTX-24-10G-W-SFPP | 7.5E-3                                             | Not supported                                  |
| MIC6-100G-CFP2       | 1.14E-5                                            | 1.0E-06                                        |
| MPC5E                | 1.14E-5                                            | 1.0E-06                                        |

To configure the threshold that clears the signal degrade alarm, include the `ber-threshold-clear` statement at the same hierarchy level. To configure the time interval during which the BER must stay above or below the configured thresholds to raise or clear the alarm, include the `interval` statement at the same hierarchy level.

**NOTE:** For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms \* 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

## Options

**value**—BER threshold for signal degradation in scientific notation. Both the mantissa and exponent are configurable. Enter the value in the format  $xE-n$ , where  $x$  is the mantissa and  $n$  is the exponent. For example, 4.5E-3.

- **Range:** The mantissa must be a decimal number. There is no limit on the number of digits before or after the decimal point.

The exponent must be an integer from 0 through 9.

- **Default:** See [Table 127 on page 702](#).

**NOTE:** In Junos OS Release 13.2R1, only the exponent is valid input for the BER threshold value, the mantissa is not configurable. The BER threshold value is  $1.0E-n$  where  $n > 0$ , and the valid range of  $n$  is from 1 through 10.

**BEST PRACTICE:** To enable proactive protection before packet loss occurs, set the **ber-threshold-signal-degrade** *value* below the FEC limit. For the FEC limits, see the table describing the signal degrade and clear thresholds after configuration in "[Understanding Pre-FEC BER Monitoring and BER Thresholds](#)" on page 543.

## 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 Release 13.2.

## RELATED DOCUMENTATION

[Understanding Pre-FEC BER Monitoring and BER Thresholds](#) | 543

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 459

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 561

## bypass

### IN THIS SECTION

- [Syntax](#) | 705
- [Hierarchy Level](#) | 705
- [Description](#) | 705

- Default | 705
- Options | 705
- Required Privilege Level | 705
- Release Information | 706

## Syntax

```
(bypass | no-bypass);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

## Description

Pass or do not pass the delay measurement (DM) value through a node.

## Default

If you omit the bypass statement, the default behavior is to disable ODU delay management options.

By default, do not pass the DM value through a node.

## Options

- |                  |                                          |
|------------------|------------------------------------------|
| <b>bypass</b>    | Pass the DM value through a node.        |
| <b>no-bypass</b> | Do not pass the DM value through a node. |

## 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 Release 13.2.

### RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## bytes (otn-options)

### IN THIS SECTION

- [Syntax | 706](#)
- [Hierarchy Level | 706](#)
- [Description | 706](#)
- [Options | 707](#)
- [Required Privilege Level | 707](#)
- [Release Information | 707](#)

### Syntax

```
bytes transmit-payload-type value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Description

Specify the transmit payload type on OTN header bytes.

## Options

*value*—Transmit payload type.

- **Range:** 0 through 255 bytes

## 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 Release 13.2.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## cfp-to-et

### IN THIS SECTION

- [Syntax | 708](#)
- [Hierarchy Level | 708](#)
- [Description | 708](#)
- [Required Privilege Level | 708](#)
- [Release Information | 708](#)

## Syntax

```
cfp-to-et;
```

## Hierarchy Level

```
[edit chassis fpc slot]
```

## Description

Make the interface et-0/1/0 (on the QSFP28 port) available for use. After you configure the **set chassis fpc 0 cfp-to-et** command and commit the configuration, you need to restart the FPC by executing the **restart chassis-control** command. After the FPC comes online, interface et-0/1/0 is created and et-0/2/1(on the CFP2 port) is deleted.

**NOTE:** Before executing this command, plan to handle disruption of services.

## 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 Release 19.2R1-S1.

## RELATED DOCUMENTATION

| [Interface Naming Conventions](#) | 394

## fec

### IN THIS SECTION

- [Syntax \(M Series, MX Series, PTX Series\) | 709](#)
- [Syntax \(ACX6360\) | 709](#)
- [Hierarchy Level \(M Series, MX Series, PTX Series\) | 709](#)
- [Hierarchy Level \(ACX6360, ACX5448-D\) | 709](#)
- [Description | 710](#)
- [Default | 710](#)
- [Options | 710](#)
- [Required Privilege Level | 711](#)
- [Release Information | 711](#)

### Syntax (M Series, MX Series, PTX Series)

```
fec (efec | gfec | gfec-sdfec |hgfec | sd-fec | ufec | none);
```

### Syntax (ACX6360)

```
fec (sdfec | sdfec15 | none);
```

### Hierarchy Level (M Series, MX Series, PTX Series)

```
[edit interfaces interface-name otn-options]
```

### Hierarchy Level (ACX6360, ACX5448-D)

```
[edit interfaces interface-name optics-options]
```

## Description

Enable forward error correction (FEC) mode.

## Default

If you do not specify a mode, the default mode is **gfec**. On PTX Series routers with P1-PTX-2-100G-WDM, the default value is **gfec-sdfec**. On PTX Series routers with PTX-5-100G-WDM and on MX Series routers with MIC3-100G-DWDM, the default value is **sdfec**.

## Options

**efec**—(M Series, MX Series routers and PTX Series routers only) G.975.1 I.4 enhanced forward error correction (EFEC) is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

**gfec**—(M series, MX Series routers and PTX Series routers only) G.709 generic forward error correction (GFEC) mode is configured to detect and correct bit errors.

**gfec-sdfec**—(PTX Series routers only) GFEC and soft-decision forward error correction (SD-FEC) modes are configured to detect and correct bit errors.

**hgfec**—(MX Series routers only) High gain forward error correction mode is configured to detect and correct bit errors.

**sdfec**—(MX Series routers, PTX Series routers, and ACX6360 routers only) Sky-Compatible Soft-decision forward error correction mode is configured to detect and correct bit errors.

**sdfec15**—(ACX6360 routers only) Soft Decision Forward Error Correction with 15 percent overhead is configured to detect and correct bit errors.

**none**—(M Series and MX Series routers only) FEC mode is not configured.

**NOTE:** On MX Series routers with MIC3-100G-DWDM and PTX Series routers with PTX-5-100G-WDM, none option is not supported. The **fec** mode must be enabled on the MIC3-100G-DWDM MIC and the PTX-5-100G-WDM PIC.

**ufec**—(MX Series routers and PTX Series routers only) G.975.1 I.7 Ultra Forward Error Correction (UFEC) mode is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

## 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 Release 9.4.

Statement and **gfec-sdfec** option introduced in Junos OS Release 13.2 for PTX Series routers. with P1-PTX-2-100G-WDM PIC.

Options **efec**, **gfec**, and **ufec** introduced in Junos OS Release 13.3 for MX Series routers. with MPC5E-100G10G, MPC5E-40G10G, MIC6-10G-OTN, and MIC6-100G-CFP2.

Options **efec**, **gfec**, and **ufec** introduced in Junos OS Release 14.1 for PTX Series routers. with P1-PTX-24-10G-W-SFPP.

Option **hgfec** introduced in Junos OS Release 15.1F5 for MX Series Routers with MIC3-100G-DWDM MIC.

Option **sdfec** introduced in Junos OS Release 15.1F5 for MX Series Routers with MIC3-100G-DWDM MIC.

Option **sdfec** introduced in Junos OS Release 15.1F6 for PTX Series Routers with PTX-5-100G-WDM PIC.

## RELATED DOCUMENTATION

---

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

---

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

---

[Supported Forward Error Correction Modes on MX Series Routers | 549](#)

---

[Supported Forward Error Correction Modes on PTX Series Routers | 549](#)

---

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 543](#)

---

[Supported Forward Error Correction Modes on ACX6360 Router | 550](#)

## fec

### IN THIS SECTION

- [Syntax | 712](#)
- [Description | 712](#)
- [Default | 712](#)
- [Options | 712](#)
- [Required Privilege Level | 713](#)
- [Release Information | 713](#)

### Syntax

```
fec (efec | gfec | gfec-sdfec | hgfec | sd-fec | ufec | none);
```

### Description

Enable forward error correction (FEC) mode.

### Default

The default value is **gfec**.

### Options

**efec**—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

**gfec**—Generic forward error correction (GFEC) mode is configured to detect and correct bit errors.

**gfec-sdfec**—GFEC and soft-decision forward error correction (SD-FEC) modes are configured to detect and correct bit errors.

**hgfec**—High gain forward error correction mode is configured to detect and correct bit errors.

**sdfec**—Soft-decision forward error correction mode is configured to detect and correct bit errors.

none—FEC mode is not configured.

ufec—Ultra Forward Error Correction (UFEC) mode is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

### 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 Release 17.2R1.

### RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 543](#)

## fec (gigether)

### IN THIS SECTION

- [Syntax | 714](#)
- [Hierarchy Level | 714](#)
- [Description | 714](#)
- [Default | 715](#)
- [Options | 715](#)
- [Required Privilege Level | 715](#)
- [Release Information | 715](#)

## Syntax

```
fec (fec91 | fec74 | none)
```

## Hierarchy Level

```
[edit interfaces interface-name together-options]
```

## Description

(ACX6360-OX Transponders, ACX6360 routers, MX Series Routers with MPC7E, MPC8E, and MPC9E, MX10003 Router with MX10003 MPC, MX204 Router, PTX1000, FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000) Enable or disable RS-FEC (Reed-Solomon Forward Error Correction) for a 100-Gigabit Ethernet interface. By default, the Junos OS software enables or disables forward error correction based on the plugged-in optics. For instance, Junos OS software enables RS-FEC for 100G SR4 optics and disables RS-FEC for 100G LR4 optics.

This statement allows you to override the default behavior and explicitly enable or disable RS-FEC. For instance, you can extend the reach of 100G LR4 optics when you explicitly enable RS-FEC for the optics. RS-FEC is compliant with IEEE 802.3-2015 Clause 91.

Once you enable or disable RS-FEC using this statement, this behavior applies to any 100-Gigabit Ethernet optical transceiver installed in the port associated with the interface.

You can configure forward error correction (FEC) clauses CL74 on 25-Gigabit and 50-Gigabit interfaces, and CL91 on 100-Gigabit interfaces. Since the FEC clauses are applied by default on these interfaces, you must disable the FEC clauses if you do not want to apply them.

**NOTE:** FPC-PTX-P1-A and FPC2-PTX-P1A on PTX5000 routers do not support RS-FEC. FPC3-SFF-PTX-1H and FP3-SFF-PTX-1T with PE-10-U-QSFP28 PIC and LR4 optics on PTX3000 and PTX5000 routers supports RS-FEC only on port 2. For PE-10-U-QSFP28 with LR4 optics, RS-FEC is the default FEC mode on port 2 and NONE is the default FEC mode on ports 0,1,3 through 9. For PE-10-U-QSFP28 with SR4 optics, RS-FEC is enabled by default on all ports. Do not modify the FEC mode on any port irrespective of the optics installed.

To check the FEC status, use the **show interfaces *interface-name*** command.

## Default

Junos OS software automatically enables or disables RS-FEC based on the type of pluggable optics used.

## Options

**fec91** Enables RS-FEC. RS-FEC is compliant with IEEE 802.3-2015 Clause 91.

**fec74** Enables RS-FEC. RS-FEC is compliant with IEEE 802.3-2015 Clause 74.

**none** Disables RS-FEC.

## 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 Release 16.1R1

## RELATED DOCUMENTATION

[Determining Transceiver Support for the PTX1000](#)

*Transponder Mode on the ACX6360*

[Understanding the features of ACX6360 | 466](#)

## fixed-stuff-bytes

### IN THIS SECTION

- [Syntax | 716](#)
- [Hierarchy Level | 716](#)
- [Description | 716](#)

- [Default | 716](#)
- [Options | 716](#)
- [Required Privilege Level | 716](#)
- [Release Information | 717](#)

## Syntax

```
(fixed-stuff-bytes | no-fixed-stuff-bytes);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

## Description

Enable or disable fixed stuff bytes.

## Default

By default, no fixed stuff bytes are set.

## Options

|                             |                                    |
|-----------------------------|------------------------------------|
| <b>fixed-stuff-bytes</b>    | Fixed stuff bytes 11.0957 Gbps.    |
| <b>no-fixed-stuff-bytes</b> | No fixed stuff bytes 11.0491 Gbps. |

## 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 Release 9.4.

### RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## high-polarization

### IN THIS SECTION

- [Syntax | 717](#)
- [Hierarchy Level | 717](#)
- [Description | 718](#)
- [Default | 718](#)
- [Required Privilege Level | 718](#)
- [Release Information | 718](#)

### Syntax

```
high-polarization;
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

## Description

Enable the physical port to rapidly track the state of polarization changes. Enabling this statement reduces the optical signal to noise ratio (OSNR) by few tenths of dB.

## Default

By default, the **high-polarization** statement is disabled.

## Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 18.2R1 for ACX6360 routers.

## RELATED DOCUMENTATION

| [Supported Optics Options on ACX6360 and ACX5448-D Routers](#) | 533

## interval

### IN THIS SECTION

- [Syntax](#) | 719
- [Hierarchy Level](#) | 719
- [Description](#) | 719
- [Options](#) | 719
- [Required Privilege Level](#) | 720
- [Release Information](#) | 720

## Syntax

```
interval value;
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

## Description

Specify the interval for which the BER must stay above the signal degradation threshold—as configured in the `ber-threshold-signal-degrade value` statement—for the alarm to be raised. After an alarm is raised, if the BER returns below the clear threshold—as configured in the `ber-threshold-clear value` statement—for the specified interval, the alarm is cleared.

**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.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options signal-degrade]` hierarchy level, then the thresholds are calculated using the pre-forward error correction (pre-FEC) BER (the BER before FEC correction). These thresholds are used for pre-FEC BER monitoring. See "[Understanding Pre-FEC BER Monitoring and BER Thresholds](#)" on page 543 for more information about pre-FEC BER monitoring and determining BER threshold settings.

If you configure the BER thresholds at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level, then the thresholds are calculated using the post-FEC BER (the BER after FEC correction). This BER is referred to as the optical channel data unit (ODU) BER.

**NOTE:** You can configure ODU BER thresholds only at the `[edit interfaces interface-name otn-options odu-signal-degrade]` hierarchy level on the P2-100GE-OTN PIC.

## Options

*value*—Time interval in milliseconds.

**NOTE:** For the P1-PTX-2-100G-WDM PIC, the BER must stay above the signal degradation threshold for ten consecutive intervals for the alarm to be raised and the BER must stay below the clear threshold for ten consecutive intervals for the alarm to be cleared. For example, if the interval is configured as 10 ms, then the BER must stay above the signal degradation threshold for 100 ms (10 ms \* 10 intervals) for the alarm to be raised, or below the clear threshold for 100 ms for the alarm to be cleared.

**NOTE:** For P1-PTX-24-10G-W-SFPP PIC and P2-100GE-OTN PIC, when the router cannot configure BER with the given interval, it selects an optimum interval that is supported for the given BER configuration. If the router is still not able to support the configuration (for example, with a wider gap between the degrade set and clear values), the default values are used and a log is generated.

For the P2-10G-40G-QSFPP PIC, the time interval is supported in multiples of 100 ms. For example, when you configure the interval as 10 ms, then it is rounded off to the nearest multiple of 100 ms.

- **Range:** 1 ms through 1000 ms.
- **Default:** 100 ms.

**NOTE:** For the P2-100GE-OTN PIC, the default value is 10 ms.

## 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 Release 13.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## is-ma

### IN THIS SECTION

- [Syntax | 721](#)
- [Hierarchy Level | 721](#)
- [Description | 721](#)
- [Default | 721](#)
- [Options | 722](#)
- [Required Privilege Level | 722](#)
- [Release Information | 722](#)

### Syntax

```
(is-ma | no-is-ma);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
[edit dynamic-profiles name interfaces name optics-options],
[edit dynamic-profiles name logical-systems name interfaces name
optics-options],
[edit dynamic-profiles name interfaces name optics-options],
[edit interfaces name optics-options]
```

### Description

Specify whether masked alarms are enabled or disabled.

### Default

If you omit the is-ma statement, masked alarms are disabled.

## Options

|                 |                              |
|-----------------|------------------------------|
| <b>is-ma</b>    | Enable masked alarms.        |
| <b>no-is-ma</b> | Do not enable masked alarms. |

## 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 Release 13.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## laser-enable

### IN THIS SECTION

- [Syntax | 723](#)
- [Hierarchy Level | 723](#)
- [Description | 723](#)
- [Default | 723](#)
- [Options | 723](#)
- [Required Privilege Level | 723](#)
- [Release Information | 723](#)

## Syntax

```
(laser-enable | no-laser-enable);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify whether lasers are enabled or disabled.

## Default

If you omit the laser-enable statement, lasers are disabled.

## Options

|                        |                       |
|------------------------|-----------------------|
| <b>laser-enable</b>    | Enable lasers.        |
| <b>no-laser-enable</b> | Do not enable lasers. |

## 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 Release 9.4.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## line-loopback

### IN THIS SECTION

- [Syntax | 724](#)
- [Hierarchy Level | 724](#)
- [Description | 724](#)
- [Default | 724](#)
- [Options | 724](#)
- [Required Privilege Level | 725](#)
- [Release Information | 725](#)

### Syntax

```
(line-loopback-enable | no-line-loopback);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Description

Specify whether line-loopback is enabled or disabled.

### Default

If you omit the line-loopback-enable statement, line-loopback is disabled.

### Options

|                             |                        |
|-----------------------------|------------------------|
| <b>line-loopback-enable</b> | Enable line-loopback.  |
| <b>no-line-loopback</b>     | Disable line-loopback. |

## 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 Release 9.4.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## local-loopback

### IN THIS SECTION

- [Syntax | 725](#)
- [Hierarchy Level | 726](#)
- [Description | 726](#)
- [Default | 726](#)
- [Options | 726](#)
- [Required Privilege Level | 726](#)
- [Release Information | 726](#)

## Syntax

```
(local-loopback-enable | no-local-loopback);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify whether local-loopback is enabled or disabled.

## Default

If you omit the local-loopback-enable statement, local-loopback is disabled.

## Options

|                              |                         |
|------------------------------|-------------------------|
| <b>local-loopback-enable</b> | Enable local-loopback.  |
| <b>no-local-loopback</b>     | Disable local-loopback. |

## 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 Release 13.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## monitor-end-point

### IN THIS SECTION

- [Syntax | 727](#)
- [Hierarchy Level | 727](#)
- [Description | 727](#)
- [Default | 727](#)
- [Options | 727](#)
- [Required Privilege Level | 728](#)
- [Release Information | 728](#)

### Syntax

```
(monitor-end-point | no-monitor-end-point);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Description

Originate or do not originate the connection monitor end point.

### Default

By default, do not originate the connection monitor end point.

### Options

|                             |                                                    |
|-----------------------------|----------------------------------------------------|
| <b>monitor-end-point</b>    | Originate the connection monitor end point.        |
| <b>no-monitor-end-point</b> | Do not originate the connection monitor end point. |

## 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 Release 13.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## no-odu-backward-frr-enable

### IN THIS SECTION

- [Syntax | 728](#)
- [Hierarchy Level | 729](#)
- [Description | 729](#)
- [Default | 729](#)
- [Required Privilege Level | 729](#)
- [Release Information | 729](#)

## Syntax

```
no-odu-backward-frr-enable;
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

## Description

Disable preemptive fast reroute (FRR) ODU backward FRR insertion.

## Default

By default, FRR ODU backward FRR insertion is disabled.

## 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 Release 14.1R2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## no-odu-signal-degrade-monitor-enable

### IN THIS SECTION

- [Syntax | 730](#)
- [Hierarchy Level | 730](#)
- [Description | 730](#)

- Default | 730
- Required Privilege Level | 730
- Release Information | 730

## Syntax

```
no-odu-signal-degrade-monitor-enable;
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

## Description

Disable monitoring of signal degradation of ODU BER in the received OTN frames.

## Default

By default, FRR signal degrade monitoring disabled.

## 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 Release 14.1R2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## number-of-frames

### IN THIS SECTION

- [Syntax | 731](#)
- [Hierarchy Level | 731](#)
- [Description | 731](#)
- [Options | 731](#)
- [Required Privilege Level | 731](#)
- [Release Information | 732](#)

### Syntax

```
number-of-frames value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Description

Specify the number of consequent frames to declare a delay measurement (DM) session completed.

### Options

*value*—Number of consequent frames to declare DM completed.

- **Range:** 0 through 255 frames.

### 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 Release 13.2.

### RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## oc192

### IN THIS SECTION

- [Syntax | 732](#)
- [Hierarchy Level | 732](#)
- [Description | 732](#)
- [Options | 733](#)
- [Required Privilege Level | 733](#)
- [Release Information | 733](#)

### Syntax

```
oc192;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

### Description

Set the line rate or speed of the OTN signal to optical channel transport unit 2 (OTU2).

## Options

**oc192** OTU2 line rate or 10 Gbps

## 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 Release 13.3 for MX Series routers.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## odu-delay-management

### IN THIS SECTION

- [Syntax | 734](#)
- [Hierarchy Level | 734](#)
- [Description | 734](#)
- [Default | 734](#)
- [Options | 734](#)
- [Required Privilege Level | 734](#)
- [Release Information | 734](#)

## Syntax

```
odu-delay-management {
 (bypass | no-bypass);
 (monitor-end-point | no-monitor-end-point);
 number-of-frames value;
 (no-start-measurement | start-measurement);
}
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify Optical Channel Data Unit (ODU) delay management options.

## Default

If you omit the `odu-delay-management` statement, the ODU delay management options are disabled.

## Options

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

## Required Privilege Level

`interface`—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 13.2.

## RELATED DOCUMENTATION

| [100-Gigabit Ethernet OTN Options Configuration Overview](#) | 459

## odu-backward-frr-enable

### IN THIS SECTION

- [Syntax | 735](#)
- [Hierarchy Level | 735](#)
- [Description | 735](#)
- [Default | 735](#)
- [Required Privilege Level | 736](#)
- [Release Information | 736](#)

### Syntax

```
odu-backward-frr-enable;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Description

Insert the ODU status into the transmitted OTN frames and monitor the received OTN frames for the ODU BER status.

### Default

By default, FRR ODU backward FRR insertion is disabled.

## 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 Release 14.1R2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## odu-signal-degrade

### IN THIS SECTION

- [Syntax | 736](#)
- [Hierarchy Level | 737](#)
- [Description | 737](#)
- [Default | 737](#)
- [Options | 737](#)
- [Required Privilege Level | 737](#)
- [Release Information | 738](#)

## Syntax

```
odu-signal-degrade {
 ber-threshold-clear;
 ber-threshold-signal-degrade;
```

```

 interval
}

```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify optical channel data unit (ODU) signal degradation threshold-related values.

## Default

If you omit the **odu-signal-degrade** statement, the default threshold values are used.

The following are the default threshold values for optical channel data unit (ODU) signal degradation for the P2-100GE-OTN PIC:

- `ber-threshold-clear`—1E-09
- `ber-threshold-signal-degrade`—1E-06
- `interval`—10 ms

The following are the default threshold values for optical channel data unit (ODU) signal degradation for the MPC5E and the MIC6-100G-CFP2 MIC:

- `ber-threshold-clear`—1.14E-5
- `ber-threshold-signal-degrade`—1.0E-6
- `interval`—10 ms

## Options

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

## Required Privilege Level

`interface`—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 14.1R2.

### RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## odu-signal-degrade-monitor-enable

### IN THIS SECTION

- [Syntax | 738](#)
- [Hierarchy Level | 738](#)
- [Description | 739](#)
- [Default | 739](#)
- [Required Privilege Level | 739](#)
- [Release Information | 739](#)

### Syntax

```
odu-signal-degrade-monitor-enable;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

## Description

Enable monitoring of signal degradation of ODU BER in the received OTN frames.

## Default

By default, FRR signal degrade monitoring disabled.

## 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 Release 14.1R2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## odu-ttim-action-enable

### IN THIS SECTION

- [Syntax | 740](#)
- [Hierarchy Level | 740](#)
- [Description | 740](#)
- [Default | 740](#)
- [Options | 740](#)
- [Required Privilege Level | 740](#)
- [Release Information | 740](#)

## Syntax

```
(odu-ttim-action-enable | no-odu-ttim-action-enable);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify whether consequent action for Optical Channel Data Unit (ODU) TTIM is enabled or disabled.

## Default

If you omit the `odu-ttim-action-enable` statement, consequent action for ODU TTIM is disabled.

## Options

|                                        |                                         |
|----------------------------------------|-----------------------------------------|
| <code>odu-ttim-action-enable</code>    | Enable consequent action for ODU TTIM.  |
| <code>no-odu-ttim-action-enable</code> | Disable consequent action for ODU TTIM. |

## 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 Release 13.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## otu-ttim-action-enable

### IN THIS SECTION

- [Syntax | 741](#)
- [Hierarchy Level | 741](#)
- [Description | 741](#)
- [Default | 741](#)
- [Options | 741](#)
- [Required Privilege Level | 742](#)
- [Release Information | 742](#)

### Syntax

```
(otu-ttim-action-enable | no-otu-ttim-action-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Description

Specify whether consequent action for Optical Channel Transport Unit (OTU) TTIM is enabled or disabled.

### Default

If you omit the `otu-ttim-action-enable` statement, consequent action for OTU TTIM is disabled.

### Options

|                               |                                        |
|-------------------------------|----------------------------------------|
| <b>otu-ttim-action-enable</b> | Enable consequent action for OTU TTIM. |
|-------------------------------|----------------------------------------|

**no-otu-ttim-action-enable** Disable consequent action for OTU TTIM.

### 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 Release 13.2.

### RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## otu4

### IN THIS SECTION

- [Syntax | 742](#)
- [Hierarchy Level | 743](#)
- [Description | 743](#)
- [Default | 743](#)
- [Options | 743](#)
- [Required Privilege Level | 743](#)
- [Release Information | 743](#)

### Syntax

```
otu4;
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

## Description

Sets the line rate or speed of the OTN signal to optical channel transport unit 4 (OTU4).

## Default

By default, the rate is OTU4 on PTX Series routers.

## Options

**otu4** OTU4 line rate or 100 Gbps

## 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 Release 13.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## pass-through

### IN THIS SECTION

- [Syntax | 744](#)
- [Hierarchy Level | 744](#)
- [Description | 744](#)
- [Default | 744](#)
- [Options | 744](#)
- [Required Privilege Level | 745](#)
- [Release Information | 745](#)

### Syntax

```
(pass-through | no-pass-through);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

### Description

Enable or disable OTN pass-through mode.

### Default

By default, OTN pass-through mode is disabled.

### Options

- |                        |                                      |
|------------------------|--------------------------------------|
| <b>no-pass-through</b> | Do not enable OTN pass-through mode. |
| <b>pass-through</b>    | Enable OTN pass-through mode.        |

## 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 Release 9.4.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## prbs

### IN THIS SECTION

- [Syntax | 745](#)
- [Hierarchy Level | 746](#)
- [Description | 746](#)
- [Default | 746](#)
- [Options | 746](#)
- [Required Privilege Level | 746](#)
- [Release Information | 746](#)

## Syntax

```
(prbs | no-prbs);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify whether OTN payload Pseudo-Random Binary Sequence (PBRs) is enabled or disabled.

## Default

By default, OTN payload prbs is disabled.

## Options

|                |                           |
|----------------|---------------------------|
| <b>prbs</b>    | Enable OTN payload PBRs.  |
| <b>no-prbs</b> | Disable OTN payload PBRs. |

## 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 Release 13.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## preemptive-fast-reroute

### IN THIS SECTION

- [Syntax | 747](#)
- [Hierarchy Level | 747](#)
- [Description | 747](#)
- [Default | 747](#)
- [Options | 748](#)
- [Required Privilege Level | 748](#)
- [Release Information | 748](#)

### Syntax

```
preemptive-fast-reroute {
 (backward-frr-enable | no-backward-frr-enable);
 (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
 (odu-backward-frr-enable | no-odu-backward-frr-enable);
 (odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable);
}
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Description

Enable or disable preemptive fast reroute options.

### Default

By default, backward fast reroute insertion and signal degradation monitoring are disabled.

## Options

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 13.2.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## rate

### IN THIS SECTION

- [Syntax | 749](#)
- [Hierarchy Level | 749](#)
- [Description | 749](#)
- [Options | 749](#)
- [Required Privilege Level | 749](#)
- [Release Information | 749](#)

## Syntax

```
rate {
 (fixed-stuff-bytes | no-fixed-stuff-bytes);
 otu4; oc192;
 (pass-through | no-pass-through);
}
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify the line rate or speed of the OTN signals.

## Options

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.4.

Statement and **otu4** option introduced in Junos OS Release 13.2 for PTX Series routers.

Option **oc192** introduced in Junos OS Release 13.3 for MX Series routers.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

## remote-loop-enable

### IN THIS SECTION

- [Syntax | 750](#)
- [Hierarchy Level | 750](#)
- [Description | 750](#)
- [Default | 751](#)
- [Options | 751](#)
- [Required Privilege Level | 751](#)
- [Release Information | 751](#)

### Syntax

```
(remote-loop-enable | no-remote-loop-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Description

Enable the remote interface to loop back the delay measurement pattern to the local interface. Delay is measured by transmitting a known pattern (delay measurement pattern) in a selected bit of the delay measurement (**DM**) field and measuring the number of frames that are missed when the delay measurement pattern is received at the transmitting end (local interface).

**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.

## Default

Delay measurement is disabled by default.

## Options

**remote-loop-enable** Enables loopback of the delay measurement pattern at the remote interface.

**no-remote-loop-enable** Disables loopback of the delay measurement pattern at the remote interface.

## 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 Release 17.1.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit DWDM OTN MIC with CFP2-ACO](#)

[100-Gigabit DWDM OTN PIC with CFP2-ACO \(PTX Series\)](#)

[Configuring OTN Interfaces on MIC3-100G-DWDM MIC | 569](#)

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 575](#)

[Disabling ODU Path Delay Measurement | 591](#)

[Enabling ODU Path Delay Measurement | 589](#)

[Understanding ODU Path Delay Measurement | 587](#)

[Understanding the MIC3-100G-DWDM MIC | 474](#)

## signal-degrade

### IN THIS SECTION

- [Syntax | 752](#)
- [Hierarchy Level | 752](#)
- [Description | 752](#)
- [Default | 753](#)
- [Options | 753](#)
- [Required Privilege Level | 753](#)
- [Release Information | 753](#)

### Syntax

```
signal-degrade {
 ber-threshold-clear value;
 ber-threshold-signal-degrade value;
 interval value;
}
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Description

Specify bit error rate (BER) signal degradation thresholds and time interval for raising and clearing alarms for optical transport network (OTN) links.

## Default

If you omit the **signal-degrade** statement, the default threshold values are used.

## Options

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 13.2.

## RELATED DOCUMENTATION

---

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

---

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## signal-degrade-monitor-enable

### IN THIS SECTION

- [Syntax | 754](#)
- [Hierarchy Level | 754](#)
- [Description | 754](#)
- [Default | 755](#)
- [Options | 755](#)
- [Required Privilege Level | 755](#)

## Syntax

```
(signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

## Description

Enable or disable pre-forward error correction (FEC) bit error rate (BER) monitoring.

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 local pre-FEC BER status into transmitted OTN 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. To configure backward fast reroute, include the `backward-frr-enable` statement at the same hierarchy level.

**NOTE:** When you configure pre-FEC BER signal degrade monitoring, we recommend that you configure both the `signal-degrade-monitor-enable` and `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. Include the `ber-threshold-signal-degrade value`, `ber-threshold-clear value`, and `interval value` statements at the `[edit interfaces interface-name otn-options signal-degrade]` hierarchy level to configure the BER thresholds and time interval. See "[Understanding Pre-FEC BER](#)"

[Monitoring and BER Thresholds" on page 543](#) for more information about pre-FEC BER monitoring and determining BER threshold settings.

## Default

By default, pre-FEC BER signal degrade monitoring is disabled.

## Options

**signal-degrade-monitor-enable**—Enable pre-FEC BER signal degrade monitoring.

**no-signal-degrade-monitor-enable**—Do not enable pre-FEC BER signal degrade monitoring.

## 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 Release 13.2.

## RELATED DOCUMENTATION

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 543](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## start-measurement

### IN THIS SECTION

- [Syntax | 756](#)
- [Hierarchy Level | 756](#)
- [Description | 756](#)

- [Default | 756](#)
- [Options | 756](#)
- [Required Privilege Level | 756](#)
- [Release Information | 757](#)

## Syntax

```
(no-start-measurement | start-measurement);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

## Description

Start or do not start a delay measurement (DM) session.

## Default

By default, do not start a DM session.

## Options

|                             |                            |
|-----------------------------|----------------------------|
| <b>no-start-measurement</b> | Do not start a DM session. |
| <b>start-measurement</b>    | Start a DM session.        |

## 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 Release 13.2.

### RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## tca

### IN THIS SECTION

- [Syntax | 757](#)
- [Hierarchy Level | 758](#)
- [Description | 758](#)
- [Default | 758](#)
- [Options | 758](#)
- [Required Privilege Level | 760](#)
- [Release Information | 760](#)

## Syntax

```
tca tca-identifier (enable-tca | no-enable-tca) (threshold number |
threshold-24hrs number)
```

## Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

```
[edit interfaces interface-name otn-options]
```

## Description

TCAs can give the management system an early indication as to the state of the associated entity when it crosses a certain threshold. TCAs can be set for both minimum and maximum values for gauges and only maximum values for counters. The timely detection of TCAs is essential to proactively manage the interface. TCAs are not an indication of a fault, but rather an indication that the entity may be close to a fault. You can choose which TCAs you want monitored by enabling the TCA. You can either keep the default threshold settings or change the settings.

Enable threshold crossing alerts (TCAs) for the following:

- Optical channel data unit (ODU)
- Optical channel transport unit (OTU)
- Laser power
- Laser temperature

## Default

By default, TCAs are not enabled.

## Options

*tca-identifier*—At the otn-options hierarchy level, it can be one of the following:

- **odu-tca-bbe**—ODU background block error threshold-crossing defect trigger
- **odu-tca-bbe-fe**—ODU far-end background block error threshold-crossing defect trigger
- **odu-tca-es**—ODU errored seconds threshold-crossing defect trigger
- **odu-tca-es-fe**—ODU far-end errored seconds threshold-crossing defect trigger
- **odu-tca-ses**—ODU severely errored seconds threshold-crossing defect trigger

- **odu-tca-ses-fe**—ODU far-end severely errored seconds threshold-crossing defect trigger
- **odu-tca-uas**—ODU unavailable seconds threshold-crossing defect trigger
- **odu-tca-uas-fe**—ODU far-end unavailable seconds threshold-crossing defect trigger
- **otu-tca-bbe**—OTU background block error threshold-crossing defect trigger
- **otu-tca-bbe-fe**—OTU far-end background block error threshold-crossing defect trigger
- **otu-tca-es**—OTU errored seconds threshold-crossing defect trigger
- **otu-tca-es-fe**—OTU far-end errored seconds threshold-crossing defect trigger
- **otu-tca-fec-ber**—OTU forward error correction bit error rate threshold-crossing defect trigger
- **otu-tca-ses**—OTU severely errored seconds threshold-crossing defect trigger
- **otu-tca-ses-fe**—OTU far-end severely errored seconds threshold-crossing defect trigger
- **otu-tca-uas**—OTU unavailable seconds threshold-crossing defect trigger
- **otu-tca-uas-fe**—OTU far-end unavailable seconds threshold-crossing defect trigger

*tca-identifier* —At the optics-options hierarchy level, it can be one of the following:

- **carrier-frequency-offset-high-tca**—Carrier frequency high threshold setting trigger
- **carrier-frequency-offset-low-tca**—Carrier frequency low threshold setting trigger
- **fec-ber**—Optics Errored Seconds Threshold crossing defect trigger
- **fec-corrected-errors-high-tca**—FEC Corrected Error High Threshold crossing defect trigger
- **fec-uncorrected-words-high-tca**—FEC Uncorrected Words High Threshold crossing defect trigger
- **laser-frequency-error-high-tca**—Laser frequency error high TCA
- **laser-frequency-error-low-tca**—Laser frequency error low TCA
- **pam-histogram-high-tca**—PAM Histogram high TCA
- **residual-isi-high-tca**—Residual ISI high TCA
- **residual-isi-low-tca**—Residual ISI low TCA
- **rx-power-high-tca**—Rx power high threshold setting trigger
- **rx-power-low-tca**—Rx power low threshold setting trigger
- **snr-low-tca**—SNR low TCA

- **tec-current-high-tca**—TEC Current high TCA
- **tec-current-low-tca**—TEC Current low TCA
- **temperature-high-tca**—Temperature high threshold setting trigger
- **temperature-low-tca**—Temperature low threshold setting trigger
- **tx-power-high-tca**—Tx power high threshold setting trigger
- **tx-power-low-tca**—Tx power low threshold setting trigger

enable-tca | no-enable-tca—To enable or disable the threshold crossing alert.

threshold | threshold-24hrs:

- **threshold *number***—Set the 15-minute interval threshold.
- **threshold-24hrs *number***—Set the 24-hour interval threshold.

## 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 Release 14.2.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

[optics-options | 1048](#)

## transport-monitoring

### IN THIS SECTION

- [Syntax | 761](#)
- [Hierarchy Level | 761](#)
- [Description | 761](#)
- [Required Privilege Level | 762](#)
- [Release Information | 762](#)

### Syntax

```
transport-monitoring;
```

### Hierarchy Level

```
[edit interfaces]
```

### Description

Monitor the performance and state of packet transport for OTN and optics modules. The following statistics are monitored:

- Packet transport for ninety-six 15-minute intervals for the current 24 hours.
- Cumulative data of the current 24 hours.
- Cumulative data of the previous 24 hours.

If this statement is configured, transport monitoring related information is shown in the output of **show interface transport** command and corresponding MIBs are available. If this option is disabled, an error is shown in the output and corresponding MIBs are not available.

## Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 14.2.

## RELATED DOCUMENTATION

| [show interfaces transport pm | 1737](#)

## trigger

### IN THIS SECTION

- [Syntax | 762](#)
- [Hierarchy Level | 763](#)
- [Description | 763](#)
- [Default | 763](#)
- [Options | 763](#)
- [Required Privilege Level | 766](#)
- [Release Information | 766](#)

## Syntax

```
trigger trigger-identifier (hold-time hold-time-value | ignore);
```

## Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

## Description

Specify defect triggers.

## Default

By default, triggers are ignored.

## Options

*trigger-identifier*—(For M Series, MX Series, SRX Series, and T Series routers only) Trigger identifier. It can be one of the following:

- **oc-lof**—Optical channel Loss of Frame defect trigger.
- **oc-lom**—Optical channel Loss of Multiframe defect trigger.
- **oc-los**—Optical channel Loss of Signal defect trigger.
- **oc-wavelength-lock**—Optical channel Wavelength Lock defect trigger.
- **odu-ais**—Optical channel data unit (ODU) Alarm Indication Signal defect trigger.
- **odu-bbe-th**—ODU Background Block Error Threshold defect trigger.
- **odu-bdi**—ODU Backward Defect Indication defect trigger.
- **odu-bei**—(MX Series routers only) ODU Backward Error Indication defect trigger.
- **odu-es-th**—ODU Errored Seconds Threshold defect trigger.
- **odu-iae**—(MX Series routers only) ODU Incoming Alignment Error defect trigger.
- **odu-lck**—ODU Locked defect trigger.
- **odu-oci**—ODU Open Connection Indication defect trigger.
- **odu-sd**—ODU Signal Degrade defect trigger.
- **odu-ses-th**—ODU Severely Errored Seconds Threshold defect trigger.

- **odu-tca-es**—(MX Series routers only) ODU Errored Seconds Threshold crossing defect trigger.
- **odu-tca-ses**—(MX Series routers only) ODU Severely Errored Seconds Threshold crossing defect trigger.
- **odu-tca-uas**—(MX Series routers only) ODU Unavailable Seconds Threshold crossing defect trigger.
- **odu-ttim**—ODU Trail Trace Identifier Mismatch defect trigger.
- **opu-ptim**—(MX Series routers only) Payload Type Identifier Mismatch defect trigger.
- **odu-uas-th**—ODU Unavailable Seconds Threshold defect trigger.
- **opu-ptm**—Optical Channel Payload (OPU) Payload Type Mismatch defect trigger.
- **otu-ais**—Optical Channel Transport Unit (OTU) Alarm Indication Signal defect trigger.
- **otu-bbe-th**—OTU Background Block Error Threshold defect trigger.
- **otu-bdi**—OTU Backward Defect Indication defect trigger.
- **otu-es-th**—OTU Errored Seconds Threshold defect trigger.
- **otu-fec-deg**—OTU FEC Degrade defect trigger.
- **otu-fec-exe**—OTU FEC Excessive Error defect trigger.
- **otu-iae**—OTU Incoming Alignment defect trigger.
- **otu-sd**—OTU Signal Degrade defect trigger.
- **otu-ses-th**—OTU Severely Errored Seconds Threshold defect trigger.
- **otu-tca-es**—(MX Series routers only) OTU Errored Seconds Threshold crossing defect trigger.
- **otu-tca-ses**—(MX Series routers only) OTU Severely Errored Seconds Threshold crossing defect trigger.
- **otu-tca-uas**—(MX Series routers only) OTU Unavailable Seconds Threshold crossing defect trigger.
- **otu-ttim**—OTU Trail Trace Identifier Mismatch defect trigger.
- **otu-uas-th**—OTU Unavailable Seconds Threshold defect trigger.

*trigger-identifier*—(For PTX Series routers only) Trigger identifier. It can be one of the following:

- **oc-lof**—Optical channel Loss of Frame defect trigger.
- **oc-lom**—Optical channel Loss of Multiframe defect trigger.
- **oc-los**—Optical channel Loss of Signal defect trigger.

- **oc-tsf**—Optical channel TOE security functionality (TSF) defect trigger.
- **oc-wavelength-lock**—Optical channel Wavelength Lock defect trigger.
- **odu-ais**—ODU Alarm Indication Signal defect trigger.
- **odu-bdi**—ODU Backward Defect Indication defect trigger.
- **odu-bei**—ODU Backward Error Indication defect trigger.
- **odu-iae**—ODU IAE defect trigger.
- **odu-lck**—ODU Locked defect trigger.
- **odu-oci**—ODU Open Connection Indication defect trigger.
- **odu-sd**—ODU Signal Degrade defect trigger.
- **odu-tca-bbe**—ODU Background Block Error Threshold crossing defect trigger.
- **odu-tca-bbe-fe**—ODU far-end Background Block Error (BEI) Threshold crossing defect trigger.
- **odu-tca-es**—ODU Errored Seconds Threshold crossing defect trigger.
- **odu-tca-es-fe**—ODU far-end Errored Seconds Threshold crossing defect trigger.
- **odu-tca-ses**—ODU Severely Errored Seconds Threshold crossing defect trigger.
- **odu-tca-ses-fe**—ODU far-end Severely Errored Seconds Threshold crossing defect trigger.
- **odu-tca-uas**—ODU Unavailable Seconds Threshold crossing defect trigger.
- **odu-tca-uas-fe**—ODU far-end Unavailable Seconds Threshold crossing defect trigger.
- **odu-ttim**—ODU Trail Trace Identifier Mismatch defect trigger.
- **opu-ptim**—Payload Type Identifier Mismatch defect trigger.
- **otu-ais**—OTU Alarm Indication Signal defect trigger.
- **otu-bdi**—OTU Backward Defect Indication defect trigger.
- **otu-fec-deg**—OTU FEC Degrade defect trigger.
- **otu-fec-exe**—OTU FEC Excessive Error defect trigger.
- **otu-iae**—OTU Incoming Alignment defect trigger.
- **otu-sd**—OTU Signal Degrade defect trigger.
- **otu-tca-bbe**—OTU Background Block Error Threshold crossing defect trigger.

- **otu-tca-bbe-fe**—OTU far-end Background Block Error (BEI) Threshold crossing defect trigger.
- **otu-tca-es**—OTU Errored Seconds Threshold crossing defect trigger.
- **otu-tca-es-fe**—OTU far-end Errored Seconds Threshold crossing defect trigger.
- **otu-tca-ses**—OTU Severely Errored Seconds Threshold crossing defect trigger.
- **otu-tca-ses-fe**—OTU far-end Severely Errored Seconds Threshold crossing defect trigger.
- **otu-tca-uas**—OTU Unavailable Seconds Threshold crossing defect trigger.
- **otu-tca-uas-fe**—OTU far-end Unavailable Seconds Threshold crossing defect trigger.
- **otu-ttim**—OTU Trail Trace Identifier Mismatch defect trigger.

hold-time *hold-time-value*—Hold time value. It can be one of the following:

- **down**—Delay before marking interface down when defect occurs (1..65534 milliseconds).
- **up**—Delay before marking interface up when defect is absent (1..65534 milliseconds).

**NOTE:** The trigger hold time value alone does not mark an interface to be up when the defect is absent or mark an interface to be down when the defect occurs. The hold time value only impacts the alarm reporting time. To mark an interface up or down, you must also configure the physical interface hold time at the [**edit interfaces *interface-name* hierarchy level**].

## 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 Release 9.4.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## tti

### IN THIS SECTION

- [Syntax | 767](#)
- [Hierarchy Level | 767](#)
- [Description | 767](#)
- [Options | 767](#)
- [Required Privilege Level | 768](#)
- [Release Information | 768](#)

### Syntax

```
tti tti-identifier;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Description

Specify trace identifier options.

### Options

*tti-identifier*—Trace identifier. It can be one of the following:

- **odu-dapi**—Optical Channel Data Unit (ODU) Destination Access Point Identifier.
- **odu-expected-receive-dapi**—ODU Expected Receive Destination Access Point Identifier.
- **odu-expected-receive-sapi**—ODU Expected Receive Source Access Point Identifier.
- **odu-sapi**—ODU Source Access Point Identifier.

- **otu-dapi**—Optical Channel Transport Unit (OTU) Destination Access Point Identifier.
- **otu-expected-receive-dapi**—OTU Expected Receive Destination Access Point Identifier.
- **otu-expected-receive-sapi**—OTU Expected Receive Source Access Point Identifier.
- **otu-sapi**—OTU Source Access Point Identifier.

### 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 Release 9.4.

### RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 561](#)

## tx-power

### IN THIS SECTION

- [Syntax | 769](#)
- [Hierarchy Level | 769](#)
- [Description | 769](#)
- [Default | 769](#)
- [Options | 769](#)
- [Required Privilege Level | 769](#)
- [Release Information | 769](#)

## Syntax

```
tx-power dbm;
```

## Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

## Description

Transmit laser output power (dBm).

## Default

If you don't specify a value, the default transmit laser output power is -2 dBm.

## Options

*dbm*—Transmit power value.

## 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 Release 13.2.

## RELATED DOCUMENTATION

[Ethernet DWDM Interface Wavelength Overview | 552](#)

[optics-options | 1048](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

## warning

### IN THIS SECTION

- [Syntax | 770](#)
- [Hierarchy Level | 770](#)
- [Description | 770](#)
- [Options | 770](#)
- [Required Privilege Level | 771](#)
- [Release Information | 771](#)

### Syntax

```
warning low-light-warning {
 (link-down | syslog);
}
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

### Description

Specifies the action to take if the receiving optics signal is below the optics low-light warning threshold. Starting in Junos OS Release 15.1, for all QSFP-based interfaces, you need not explicitly configure the **syslog** option. The **syslog** option is enabled by default.

### Options

**link-down**—Drop the 10-Gigabit Ethernet link and marks link as down.

**syslog**—Write the optics information to the system log.

## 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 Release 10.0.

## RELATED DOCUMENTATION

[Configuring Link Down Notification for Optics Options Alarm or Warning | 214](#)

[optics-options | 1048](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

## wavelength

### IN THIS SECTION

- [Syntax | 771](#)
- [Hierarchy Level | 772](#)
- [Description | 772](#)
- [Options | 772](#)
- [Required Privilege Level | 776](#)
- [Release Information | 776](#)

## Syntax

```
wavelength nm;
```

## Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

## Description

For 10-Gigabit or 100-Gigabit Ethernet DWDM interfaces only, configure full C-band ITU-Grid tunable optics.

## Options

*nm*—Wavelength value. It can be one of the following:

**NOTE:** All values are displayed. However, if you configure a value that is not supported by the device, an error message is displayed and the device is not tuned to the specified wavelength.

- **1528.38**—1528.38 nanometers (nm), corresponds to a 50-GHz grid
- **1528.77**—1528.77 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.16**—1529.16 nm, corresponds to a 50-GHz grid
- **1529.55**—1529.55 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.94**—1529.94 nm, corresponds to a 50-GHz grid
- **1530.33**—1530.33 nm, corresponds to 50-GHz and 100-GHz grids
- **1530.72**—1530.72 nm, corresponds to a 50-GHz grid
- **1531.12**—1531.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1531.51**—1531.51 nm, corresponds to a 50-GHz grid
- **1531.90**—1531.90 nm, corresponds to 50-GHz and 100-GHz grids
- **1532.29**—1532.29 nm, corresponds to a 50-GHz grid
- **1532.68**—1532.68 nm, corresponds to 50-GHz and 100-GHz grids
- **1533.07**—1533.07 nm, corresponds to a 50-GHz grid
- **1533.47**—1533.47 nm, corresponds to 50-GHz and 100-GHz grids

- 1533.86–1533.86 nm, corresponds to a 50-GHz grid
- 1534.25–1534.25 nm, corresponds to 50-GHz and 100-GHz grids
- 1534.64–1534.64 nm, corresponds to a 50-GHz grid
- 1535.04–1535.04 nm, corresponds to 50-GHz and 100-GHz grids
- 1535.43–1535.43 nm, corresponds to a 50-GHz grid
- 1535.82–1535.82 nm, corresponds to 50-GHz and 100-GHz grids
- 1536.22–1536.22 nm, corresponds to a 50-GHz grid
- 1536.61–1536.61 nm, corresponds to 50-GHz and 100-GHz grids
- 1537.00–1537.00 nm, corresponds to a 50-GHz grid
- 1537.40–1537.40 nm, corresponds to 50-GHz and 100-GHz grids
- 1537.79–1537.79 nm, corresponds to a 50-GHz grid
- 1538.19–1538.19 nm, corresponds to 50-GHz and 100-GHz grids
- 1538.58–1538.58 nm, corresponds to a 50-GHz grid
- 1538.98–1538.98 nm, corresponds to 50-GHz and 100-GHz grids
- 1539.37–1539.37 nm, corresponds to a 50-GHz grid
- 1539.77–1539.77 nm, corresponds to 50-GHz and 100-GHz grids
- 1540.16–1540.16 nm, corresponds to a 50-GHz grid
- 1540.56–1540.56 nm, corresponds to 50-GHz and 100-GHz grids
- 1540.95–1540.95 nm, corresponds to a 50-GHz grid
- 1541.35–1541.35 nm, corresponds to 50-GHz and 100-GHz grids
- 1541.75–1541.75 nm, corresponds to a 50-GHz grid
- 1542.14–1542.14 nm, corresponds to 50-GHz and 100-GHz grids
- 1542.54–1542.54 nm, corresponds to a 50-GHz grid
- 1542.94–1542.94 nm, corresponds to 50-GHz and 100-GHz grids
- 1543.33–1543.33 nm, corresponds to a 50-GHz grid
- 1543.73–1543.73 nm, corresponds to 50-GHz and 100-GHz grids

- 1544.13–1544.13 nm, corresponds to a 50-GHz grid
- 1544.53–1544.53 nm, corresponds to 50-GHz and 100-GHz grids
- 1544.92–1544.92 nm, corresponds to a 50-GHz grid
- 1545.32–1545.32 nm, corresponds to 50-GHz and 100-GHz grids
- 1545.72–1545.72 nm, corresponds to a 50-GHz grid
- 1546.12–1546.12 nm, corresponds to 50-GHz and 100-GHz grids
- 1546.52–1546.52 nm, corresponds to a 50-GHz grid
- 1546.92–1546.92 nm, corresponds to 50-GHz and 100-GHz grids
- 1547.32–1547.32 nm, corresponds to a 50-GHz grid
- 1547.72–1547.72 nm, corresponds to 50-GHz and 100-GHz grids
- 1548.11–1548.11 nm, corresponds to a 50-GHz grid
- 1548.51–1548.51 nm, corresponds to 50-GHz and 100-GHz grids
- 1548.91–1548.91 nm, corresponds to a 50-GHz grid
- 1549.32–1549.32 nm, corresponds to 50-GHz and 100-GHz grids
- 1549.72–1549.72 nm, corresponds to a 50-GHz grid
- 1550.12–1550.12 nm, corresponds to 50-GHz and 100-GHz grids
- 1550.52–1550.52 nm, corresponds to a 50-GHz grid
- 1550.92–1550.92 nm, corresponds to 50-GHz and 100-GHz grids
- 1551.32–1551.32 nm, corresponds to a 50-GHz grid
- 1551.72–1551.72 nm, corresponds to 50-GHz and 100-GHz grids
- 1552.12–1552.12 nm, corresponds to a 50-GHz grid
- 1552.52–1552.52 nm, corresponds to 50-GHz and 100-GHz grids
- 1552.93–1552.93 nm, corresponds to a 50-GHz grid
- 1553.33–1554.33 nm, corresponds to 50-GHz and 100-GHz grids
- 1553.73–1554.73 nm, corresponds to a 50-GHz grid
- 1554.13–1554.13 nm, corresponds to 50-GHz and 100-GHz grids

- 1554.54–1554.54 nm, corresponds to a 50-GHz grid
- 1554.94–1554.94 nm, corresponds to 50-GHz and 100-GHz grids
- 1555.34–1555.34 nm, corresponds to a 50-GHz grid
- 1555.75–1555.75 nm, corresponds to 50-GHz and 100-GHz grids
- 1556.15–1556.15 nm, corresponds to a 50-GHz grid
- 1556.55–1556.55 nm, corresponds to 50-GHz and 100-GHz grids
- 1556.96–1556.96 nm, corresponds to a 50-GHz grid
- 1557.36–1557.36 nm, corresponds to 50-GHz and 100-GHz grids
- 1557.77–1557.77 nm, corresponds to a 50-GHz grid
- 1558.17–1558.17 nm, corresponds to 50-GHz and 100-GHz grids
- 1558.58–1558.58 nm, corresponds to a 50-GHz grid
- 1558.98–1558.98 nm, corresponds to 50-GHz and 100-GHz grids
- 1559.39–1559.39 nm, corresponds to a 50-GHz grid
- 1559.79–1559.79 nm, corresponds to 50-GHz and 100-GHz grids
- 1560.20–1560.20 nm, corresponds to a 50-GHz grid
- 1560.61–1560.61 nm, corresponds to 50-GHz and 100-GHz grids
- 1561.01–1561.01 nm, corresponds to a 50-GHz grid
- 1561.42–1561.42 nm, corresponds to 50-GHz and 100-GHz grids
- 1561.83–1561.83 nm, corresponds to a 50-GHz grid
- 1562.23–1562.23 nm, corresponds to 50-GHz and 100-GHz grids
- 1562.64–1562.64 nm, corresponds to a 50-GHz grid
- 1563.05–1563.05 nm, corresponds to 50-GHz and 100-GHz grids
- 1563.45–1563.45 nm, corresponds to a 50-GHz grid
- 1563.86–1563.86 nm, corresponds to 50-GHz and 100-GHz grids
- 1564.27–1564.27 nm, corresponds to a 50-GHz grid
- 1564.68–1564.68 nm, corresponds to 50-GHz and 100-GHz grids

- **1565.09**–1565.09 nm, corresponds to a 50-GHz grid
- **1565.50**–1565.50 nm, corresponds to 50-GHz and 100-GHz grids
- **1565.90**–1565.90 nm, corresponds to a 50-GHz grid
- **1566.31**–1566.31 nm, corresponds to 50-GHz and 100-GHz grids
- **1566.72**–1566.72 nm, corresponds to a 50-GHz grid
- **1567.13**–1567.13 nm, corresponds to 50-GHz and 100-GHz grids
- **1567.54**–1567.54 nm, corresponds to a 50-GHz grid
- **1567.95**–1567.95 nm, corresponds to 50-GHz and 100-GHz grids
- **1568.36**–1568.36 nm, corresponds to a 50-GHz grid
- **1568.77**–1568.77 nm, corresponds to 50-GHz and 100-GHz grids
- **Default: 1550.12**–1550.12 nm, corresponds to 50-GHz and 100-GHz grids

### Required Privilege Level

interface—To view this statement in the configuration.

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

### Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

[Ethernet DWDM Interface Wavelength Overview | 552](#)

[Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength | 553](#)

[show interfaces diagnostics optics \(Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port\) | 1368](#)

# Configuration Statements

## IN THIS CHAPTER

- `accept-source-mac` | 784
- `access-concentrator` | 786
- `account-layer2-overhead` (PIC Level) | 788
- `adaptive` | 789
- `address` | 791
- `age` | 795
- `agent-specifier` | 796
- `aggregate` (Gigabit Ethernet CoS Policer) | 798
- `aggregated-ether-options` | 800
- `alarms` | 802
- `allow-remote-loopback` | 803
- `asynchronous-notification` | 804
- `auto-negotiation` | 806
- `auto-reconnect` | 808
- `bandwidth-limit` (Policer for Gigabit Ethernet Interfaces) | 810
- `bert-algorithm` | 811
- `bert-error-rate` | 814
- `bert-period` | 816
- `bridge-domain` | 818
- `burst-size-limit` (Policer for Gigabit Ethernet Interfaces) | 820
- `centralized` | 821
- `ces-psn-channel` (tdm-options) | 823
- `cfp-to-et` | 824
- `classifier` | 826
- `client` | 827
- `community-vlans` (MX Series) | 828

- connectivity-fault-management | 830
- delay (PPPoE Service Name Tables) | 834
- destination (IPCP) | 835
- device-count | 837
- drop (PPPoE Service Name Tables) | 838
- dynamic-profile (PPPoE Service Name Tables) | 840
- egress-policer-overhead | 842
- encapsulation (Logical Interface) | 843
- encapsulation | 848
- ether-options | 856
- ethernet (Chassis) | 865
- ethernet-policer-profile | 867
- evcs | 869
- family | 870
- family | 872
- fastether-options | 878
- flow-control | 880
- fnp | 882
- force-up | 883
- forwarding-class (Gigabit Ethernet IQ Classifier) | 885
- forwarding-mode (100-Gigabit Ethernet) | 886
- forwarding-mode (PTX Series Packet Transport Routers) | 888
- frame-error | 889
- frame-period | 891
- frame-period-summary | 892
- framing (10-Gigabit Ethernet Interfaces) | 894
- ggether-options | 896
- hash-key (Chassis LAG) | 905
- hold-time up | 907
- ieee802.1p | 908
- ignore-l3-incompletes | 910
- inet (chassis) | 911

- ingress-policer-overhead | 912
- ingress-rate-limit | 915
- inline | 916
- input-policer | 918
- input-priority-map | 920
- input-three-color | 921
- input-vlan-map (Aggregated Ethernet) | 923
- interface (LLDP) | 924
- interface (OAM Link-Fault Management) | 927
- interface-group | 929
- interface-group-down | 930
- interface-none | 931
- isolated-vlan (MX Series) | 932
- iwf-params (tdm-options) | 934
- lacp (802.3ad) | 936
- lacp (Aggregated Ethernet) | 937
- lacp | 940
- lacp (Protocols) | 942
- lane | 944
- lane-all | 947
- layer2-policer | 950
- link-adjacency-loss | 952
- link-discovery | 954
- link-degrade-monitor | 955
- link-down | 958
- link-event-rate | 959
- link-fault-management | 960
- link-mode | 962
- link-protection | 964
- link-protection (non-LACP) | 966
- link-protection (Protocols LACP) | 968
- link-speed (Aggregated Ethernet) | 969

- link-speed (Aggregated SONET/SDH) | 972
- lmi (Ethernet OAM) | 973
- load-balance | 975
- load-balance-stateful (Aggregated Ethernet Interfaces) | 977
- load-type (Aggregated Ethernet Interfaces) | 979
- local-bias (ae load-balance) | 981
- logical-tunnel-options | 982
- loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | 984
- loopback (Local and Remote) | 986
- loopback-tracking | 988
- loss-priority | 989
- loopback-remote | 990
- mac | 991
- mac-address (Accept Source Mac) | 993
- mac-learn-enable | 994
- mac-validate | 996
- master-only | 997
- max-sessions (PPPoE Service Name Tables) | 998
- max-sessions-vs-a-ignore (Static and Dynamic Subscribers) | 1000
- maximum-links | 1002
- mc-ae | 1004
- minimum-bandwidth (aggregated Ethernet) | 1010
- minimum-links | 1012
- mixed-rate-mode | 1014
- mtu | 1016
- mru | 1020
- multicast-statistics | 1021
- multiservice | 1023
- negotiate-address | 1025
- negotiation-options | 1026
- no-adaptive | 1027
- no-allow-link-events | 1029

- no-auto-mdix | 1030
- no-keepalives | 1031
- no-pre-classifier | 1033
- no-send-pads-ac-info | 1034
- no-send-pads-error | 1036
- non-revertive (Chassis) | 1037
- non-revertive (Interfaces) | 1038
- number-of-ports | 1040
- number-of-sub-ports | 1042
- oam | 1044
- optics-options | 1048
- otn-options | 1051
- output-policer | 1054
- output-priority-map | 1055
- output-three-color | 1057
- output-vlan-map (Aggregated Ethernet) | 1058
- pado-advertise | 1060
- passive-monitor-mode | 1061
- payload | 1062
- pdu-interval | 1064
- pdu-threshold | 1066
- per-flow (Aggregated Ethernet Interfaces) | 1067
- periodic | 1069
- pic-mode | 1071
- policer (CFM Firewall) | 1073
- policer (CoS) | 1074
- policer (MAC) | 1076
- port-priority | 1078
- pp0 (Dynamic PPPoE) | 1079
- ppm (Ethernet Switching) | 1082
- pppoe-options | 1083
- pppoe-underlying-options (Static and Dynamic Subscribers) | 1085

- preferred-source-address | 1087
- premium (Output Priority Map) | 1089
- premium (Policer) | 1090
- protocol-down | 1091
- premium (Output Priority Map) | 1092
- premium (Policer) | 1094
- proxy | 1095
- rebalance (Aggregated Ethernet Interfaces) | 1097
- receive-options-packets | 1098
- receive-ttl-exceeded | 1100
- recovery | 1101
- remote-loopback | 1103
- restore-interval | 1104
- revertive | 1106
- routing-instance | 1107
- routing-instance (PPPoE Service Name Tables) | 1109
- rx-enable | 1110
- rx-max-duration | 1112
- sa-multicast (100-Gigabit Ethernet) | 1113
- sa-multicast (PTX Series Packet Transport Routers) | 1115
- send-critical-event | 1116
- server | 1117
- service (PPPoE) | 1119
- service-name | 1121
- service-name-table | 1122
- service-name-tables | 1124
- session-expiry (MX Series in Enhanced LAN Mode) | 1126
- sonet | 1128
- source-address-filter | 1129
- source-filtering | 1131
- speed | 1132
- speed | 1134

- speed (Ethernet) | 1137
- speed (MX Series DPC) | 1148
- speed (Gigabit Ethernet interface) | 1150
- speed (24-port and 12-port 10 Gigabit Ethernet PIC) | 1152
- static-interface | 1154
- switch-options | 1156
- switch-port | 1157
- symbol-period | 1158
- symmetric-hash | 1160
- sync-reset | 1161
- syslog (OAM Action) | 1163
- system-id | 1164
- system-priority | 1166
- system-priority | 1167
- targeted-options (Grouping Subscribers by Bandwidth Usage) | 1169
- targeted-options (Manual Targeting) | 1171
- targeted-distribution | 1173
- targeted-options | 1174
- tdm-options (Interfaces) | 1176
- terminate (PPPoE Service Name Tables) | 1178
- thresholds | 1179
- traceoptions (LLDP) | 1181
- traceoptions (Individual Interfaces) | 1185
- traceoptions (LACP) | 1194
- traceoptions (PPPoE) | 1197
- tx-duration | 1200
- tx-enable | 1202
- underlying-interface | 1203
- unit | 1205
- unnumbered-address (Dynamic Profiles) | 1215
- unnumbered-address (PPP) | 1219
- Unused | 1220

- [virtual-control-channel](#) | 1222
- [virtual-switch](#) | 1224
- [vlan-rule \(100-Gigabit Ethernet Type 4 PIC with CFP\)](#) | 1225
- [vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\)](#) | 1226

## accept-source-mac

### IN THIS SECTION

- [Syntax](#) | 784
- [Hierarchy Level](#) | 785
- [Description](#) | 785
- [Required Privilege Level](#) | 785
- [Release Information](#) | 786

### Syntax

```
accept-source-mac {
 mac-address mac-address {
 policer {
 input cos-policer-name;
 output cos-policer-name;
 }
 }
}
```

## Hierarchy Level

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

## Description

For Gigabit Ethernet intelligent queuing (IQ) interfaces only or for Gigabit Ethernet and aggregated Ethernet interfaces on switches, accept traffic from and to the specified remote media access control (MAC) address.

The **accept-source-mac** statement is equivalent to the **source-address-filter** statement, which is valid for aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only. To allow the interface to receive packets from specific MAC addresses, include the **accept-source-mac** statement.

On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

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

**NOTE:** The **policer** statement is not supported on PTX Series Packet Transport Routers.

**NOTE:** If you configure MAC filtering on the AE interface, you must configure the interface with **family ethernet-switching**.

**NOTE:** On QFX platforms, if you configure source MAC addresses for an interface using the *static-mac* or *persistent-learning* statements and later configure a different MAC address for the same interface using the **accept-source-mac** statement, the MAC addresses that you previously configured for the interface remain in the ethernet-switching table and can still be used to send packets to the interface.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[Configuring MAC Address Filtering on PTX Series Packet Transport Routers](#)

[source-filtering | 1131](#)

## access-concentrator

### IN THIS SECTION

- [Syntax | 786](#)
- [Hierarchy Level | 786](#)
- [Description | 787](#)
- [Options | 787](#)
- [Required Privilege Level | 787](#)
- [Release Information | 787](#)

## Syntax

```
access-concentrator name;
```

## Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number
family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-
```

```

number family pppoe],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-options],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-
options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number pppoe-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number pppoe-underlying-options]

```

## Description

Configure an alternative access concentrator name in the AC-NAME tag in a PPPoE control packet for use with a dynamic PPPoE subscriber interface. If you do not configure the access concentrator name, the AC-NAME tag contains the system name.

**NOTE:** The [edit ... family pppoe] hierarchies are supported only on MX Series routers with MPCs.

## Options

*name*—Name of the access concentrator.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options] and [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-underlying-options] hierarchy levels introduced in Junos OS Release 10.1.

Support at the [edit ... family pppoe] hierarchies introduced in Junos OS Release 11.2.

## RELATED DOCUMENTATION

[Identifying the Access Concentrator](#)

*[Configuring the PPPoE Family for an Underlying Interface](#)*

*[Configuring Dynamic PPPoE Subscriber Interfaces](#)*

[PPPoE Overview](#)

## account-layer2-overhead (PIC Level)

### IN THIS SECTION

- [Syntax | 788](#)
- [Hierarchy Level | 788](#)
- [Description | 788](#)
- [Required Privilege Level | 789](#)
- [Release Information | 789](#)

### Syntax

```
account-layer2-overhead;
```

### Hierarchy Level

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

### Description

Enable the automatic adjustment of Layer 2 overhead in bytes, which is the octet adjustment per packet, based on the encapsulation on the logical interface for the total octet count for ingress and egress traffic on all the interfaces in the PIC.

## 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 Release 13.2.

## RELATED DOCUMENTATION

[Accounting of the Layer 2 Overhead Attribute in Interface Statistics | 260](#)

[Configuring Layer 2 Overhead Accounting in Interface Statistics | 263](#)

[Verifying the Accounting of Layer 2 Overhead in Interface Statistics | 264](#)

## adaptive

### IN THIS SECTION

- [Syntax | 789](#)
- [Hierarchy Level | 790](#)
- [Description | 790](#)
- [Options | 790](#)
- [Required Privilege Level | 791](#)
- [Release Information | 791](#)

## Syntax

```
adaptive {
 pps;
 scan-interval multiple;
```

```

 tolerance tolerance-percentage;
}

```

## Hierarchy Level

```

[edit dynamic-profiles name interfaces name aggregated-ether-options load-
balance],
[edit dynamic-profiles name interfaces name logical-tunnel-options load-balance],
[edit dynamic-profiles name interfaces interface-range name aggregated-ether-
options load-balance],
[edit dynamic-profiles name interfaces interface-range name logical-tunnel-
options load-balance],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-
ether-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces name logical-tunnel-
options load-balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name
aggregated-ether-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name
logical-tunnel-options load-balance],
[edit interfaces name aggregated-ether-options load-balance],
[edit interfaces name logical-tunnel-options load-balance],
[edit interfaces interface-range name aggregated-ether-options load-
balance],
[edit interfaces interface-range name logical-tunnel-options load-
balance]

```

## Description

Correct a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an aggregated Ethernet bundle.

## Options

|                                          |                                                                                                                                                          |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>pps</b>                               | (PTX Series only) The type of traffic rate among the members of the AE bundle is measured packets per second. The default rate type is bytes per second. |
| <b>scan-interval<br/><i>multiple</i></b> | (PTX Series only) Scan interval, as a multiple of a 30-second interval.                                                                                  |

- **Range:** 1 through 5
- **Default:** 1

**tolerance**  
*tolerance-*  
*percentage*

(MX Series and PTX Series) Limit to the variance in the packet traffic flow to the aggregated Ethernet links in a percentage.

- **Range:** 1 through 100 percent
- **Default:** 20 percent

### 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 Release 13.2R3.

### RELATED DOCUMENTATION

*Understanding Aggregated Ethernet Load Balancing*

*Example: Configuring Aggregated Ethernet Load Balancing*

## address

### IN THIS SECTION

- [Syntax | 792](#)
- [Hierarchy Level | 793](#)
- [Description | 793](#)
- [Options | 794](#)
- [Required Privilege Level | 794](#)
- [Release Information | 794](#)

## Syntax

```

address address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 destination address;
 destination-profile name;
 eui-64;
 primary-only;
 multipoint-destination address dlci dlci-identifier;
 multipoint-destination address {
 epd-threshold cells;
 inverse-arp;
 oam-liveness {
 up-count cells;
 down-count cells;
 }
 oam-period (disable | seconds);
 shaping {
 (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak
rate sustained rate burst length);
 queue-length number;
 }
 vci vpi-identifier.vci-identifier;
 }
 primary;
 preferred;
 virtual-gateway-address
(vrrp-group | vrrp-inet6-group) group-number {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 authentication-type authentication;
 authentication-key key;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 priority-number number;
 track {
 priority-cost seconds;
 priority-hold-time interface-name {
 interface priority;
 }
 }
}

```

```

 bandwidth-threshold bits-per-second {
 priority;
 }
 }
 route ip-address/mask routing-instance instance-name priority-cost
cost;
}
virtual-address [addresses];
}
}

```

## Hierarchy Level

```

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

```

## Description

Configure the interface address.

**NOTE:** If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, and the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```

interfaces {
 xe-0/0/0 {
 unit 0 {
 family inet {
 address 192.168.1.1/8;
 }
 }
 }
 xe-0/0/1 {
 unit 0 {

```

```

family inet {
 address 192.168.1.1/8;
}
}
}

```

For more information on configuring the same address on multiple interfaces, see [Configuring the Interface Address](#).

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

## Options

*address*—Address of the interface.

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

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring the Protocol Family](#)

[family](#)

[negotiate-address](#)

*unnumbered-address (Ethernet)*

## age

### IN THIS SECTION

- [Syntax | 795](#)
- [Hierarchy Level | 795](#)
- [Description | 795](#)
- [Default | 795](#)
- [Required Privilege Level | 796](#)
- [Release Information | 796](#)

### Syntax

```
age (30m | 10m | 1m | 30s | 10s);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management linktrace]
```

### Description

Time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.

### Default

10 minutes

## 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 Release 8.5.

## RELATED DOCUMENTATION

| *Configuring Linktrace Protocol in CFM*

## agent-specifier

### IN THIS SECTION

- [Syntax | 796](#)
- [Hierarchy Level | 797](#)
- [Description | 797](#)
- [Default | 797](#)
- [Options | 797](#)
- [Required Privilege Level | 798](#)
- [Release Information | 798](#)

## Syntax

```
agent-specifier {
 aci circuit-id-string ari remote-id-string {
 drop;
 delay seconds;
 terminate;
 dynamic-profile profile-name;
```

```

 routing-instance routing-instance-name;
 static-interface interface-name;
}
}

```

## Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name]
```

## Description

Specify the action taken by the interface for the specified agent circuit identifier/agent remote identifier (ACI/ARI) pair when the interface receives a PPPoE Active Discovery Initiation (PADI) control packet that includes the vendor-specific tag with ACI/ARI pair information. You can configure an ACI/ARI pair for a named service, **empty** service, or **any** service in a PPPoE service name table. A maximum of 8000 ACI/ARI pairs are supported per PPPoE service name table. You can distribute the ACI/ARI pairs in any combination among the named, **empty**, and **any** service entries in the service name table.

You can use an asterisk (\*) as a wildcard character to match ACI/ARI pairs, the ACI alone, or the ARI alone. The asterisk can be placed only at the beginning, the end, or both the beginning and end of the identifier string. You can also specify an asterisk alone for either the ACI or the ARI. You cannot specify only an asterisk for both the ACI and the ARI. When you specify a single asterisk as the identifier, that identifier is ignored in the PADI packet.

For example, suppose you care about matching only the ACI and do not care what value the ARI has in the PADI packet, or even whether the packet contains an ARI value. In this case you can set the *remote-id-string* to a single asterisk. Then the interface ignores the ARI received in the packet and the interface takes action based only on matching the specified ACI.

## Default

The default action is terminate.

## Options

**aci *circuit-id-string***—Identifier for the agent circuit ID that corresponds to the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

**ari *remote-id-string***—Identifier for the subscriber associated with the DSLAM interface that initiated the service request. This is a string of up to 63 characters.

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

### Release Information

Statement introduced in Junos OS Release 10.0.

**drop**, **delay**, **terminate**, **dynamic-profile**, **routing-instance**, and **static-interface** options introduced in Junos OS Release 10.2.

### RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

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

## aggregate (Gigabit Ethernet CoS Policer)

### IN THIS SECTION

- [Syntax | 799](#)
- [Hierarchy Level | 799](#)
- [Description | 799](#)
- [Required Privilege Level | 799](#)
- [Release Information | 799](#)

## Syntax

```
aggregate {
 bandwidth-limit bps;
 burst-size-limit bytes;
}
```

## Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile policer cos-policer-name]
```

## Description

Define a policer to apply to nonpremium traffic.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

*premium (Hierarchical Policer)*

[ieee802.1p | 908](#)

## aggregated-ether-options

### IN THIS SECTION

- Syntax | 800
- Hierarchy Level | 801
- Description | 801
- Required Privilege Level | 802
- Release Information | 802

### Syntax

```
aggregated-ether-options {
 ethernet-switch-profile {
 ethernet-policer-profile {
 input-priority-map {
 ieee802.1p premium [values];
 }
 output-priority-map {
 classifier {
 premium {
 forwarding-class class-name {
 loss-priority (high | low);
 }
 }
 }
 }
 }
 policer cos-policer-name {
 aggregate {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 premium {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 }
 }
}
```

```

 }
 (mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
lacp {
 (active | passive);
 link-protection {
 disable;
 }
 (revertive | non-revertive);
 periodic interval;
 sync-reset
 system-priority priority;
 system-id system-id;
}
load-balance {
 local-bias;
 no-adaptive;
 per-packet;
}
local-bias;
link-speed speed;
logical-interface-chassis-redundancy;
logical-interface-fpc-redundancy;
(loopback | no-loopback);
minimum-links number;
rebalance-periodic time hour:minute <interval hours>;
source-address-filter {
 mac-address;
 (source-filtering | no-source-filtering);
}
}

```

## Hierarchy Level

```
[edit interfaces aex]
```

## Description

Configure aggregated Ethernet-specific interface properties.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Ethernet Interfaces Overview](#) | 2

## alarms

### IN THIS SECTION

- [Syntax](#) | 802
- [Hierarchy Level](#) | 802
- [Description](#) | 803
- [Required Privilege Level](#) | 803
- [Release Information](#) | 803

## Syntax

```
alarms;
```

## Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

## Description

For 10-Gigabit Ethernet DPCs, configure the DPC to drop the interface link when the receive power falls below the alarm threshold.

## 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 Release 10.1.

## RELATED DOCUMENTATION

| [Ethernet DWDM Interface Wavelength Overview](#) | 552

# allow-remote-loopback

### IN THIS SECTION

- [Syntax](#) | 803
- [Hierarchy Level](#) | 804
- [Description](#) | 804
- [Required Privilege Level](#) | 804
- [Release Information](#) | 804

## Syntax

```
allow-remote-loopback;
```

## Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name negotiation-
options]
```

## Description

Enable the remote loopback on IQ2 and IQ2-E Gigabit Ethernet interfaces, and Ethernet interfaces on the MX Series routers and EX Series switches.

## 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 Release 8.4.

## RELATED DOCUMENTATION

| *Enabling Remote Loopback Support on the Local Interface*

## asynchronous-notification

### IN THIS SECTION

- [Syntax | 805](#)
- [Hierarchy Level | 805](#)
- [Description | 805](#)
- [Default | 805](#)
- [Required Privilege Level | 805](#)
- [Release Information | 805](#)

## Syntax

```
(asynchronous-notification | no-asynchronous-notification);
```

## Hierarchy Level

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

## Description

(MX Series routers, T Series routers) For all Gigabit Ethernet interfaces (1-Gigabit, 10-Gigabit, and 100-Gigabit), configure support for notification of link down alarm generation and transfer.

(M120 and M320 routers) For all 10-Gigabit Ethernet PIC interfaces, configure support for notification of link down alarm generation and transfer.

- **asynchronous-notification**—Support notification of link down alarm generation and transfer.
- **no-asynchronous-notification**—Prohibit notification of link down alarm generation and transfer.

## Default

Support for notification of link down alarm generation and transfer is not enabled.

## 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 Release 8.3.

## RELATED DOCUMENTATION

[Gigabit Ethernet Notification of Link Down Alarm Overview | 212](#)

[Configuring Gigabit Ethernet Notification of Link Down Alarm | 213](#)

## auto-negotiation

### IN THIS SECTION

- Syntax | 806
- Hierarchy Level | 806
- Description | 806
- Default | 807
- Options | 808
- Required Privilege Level | 808
- Release Information | 808

### Syntax

```
(auto-negotiation | no-auto-negotiation) <remote-fault (local-interface-online | local-interface-offline)>;
```

### Hierarchy Level

```
[edit interfaces interface-name ether-options],
[edit interfaces interface-name gigether-options],
[edit interfaces ge-pim/0/0 switch-options switch-port port-number]
```

### Description

For Gigabit Ethernet interfaces on M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers explicitly enable autonegotiation and remote fault. (krupac)For EX Series switches, explicitly enable autonegotiation only. You cannot disable autonegotiation on ACX5448 routers for Gigabit Ethernet interfaces by using the **no-auto-negotiation** command. The **no-auto-negotiation** command is not supported on ACX5448 routers.

- **auto-negotiation**—Enables autonegotiation. This is the default. krupac

- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.

**NOTE:** On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the **no-auto-negotiation** statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

On EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.

**NOTE:** On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

**NOTE:** On ACX Series routers, except ACX5448 routers, when you configure fiber interfaces (fiber media mode) to operate at 1 Gbps, autonegotiation is enabled by default to negotiate the speed and duplex settings. You can disable autonegotiation by using the (**no-auto-negotiation**) statement, and commit the configuration. in the fiber media mode. In copper interfaces (copper media mode), autonegotiation is enabled by default. To disable autonegotiation, you need to explicitly configure the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

For SRX Series devices, when autonegotiation is disabled, you can set the **mdi-mode** to enable it in case of non-cross table.

## Default

Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

## Options

**remote-fault (local-interface-online | local-interface-offline)**—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

- **Default:** local-interface-online

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 7.6.

## RELATED DOCUMENTATION

[Gigabit Ethernet Autonegotiation Overview | 284](#)

[Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

*Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support*

## auto-reconnect

### IN THIS SECTION

- [Syntax | 809](#)
- [Hierarchy Level | 809](#)
- [Description | 809](#)
- [Options | 809](#)
- [Required Privilege Level | 809](#)
- [Release Information | 809](#)

## Syntax

```
auto-reconnect seconds;
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number pppoe-options],
[edit logical-systems logical-system-name interfaces interface-name
 unit logical-unit-number pppoe-options]
```

## Description

PPP over Ethernet interfaces, configure the amount of time to wait before reconnecting after a session has terminated.

## Options

*seconds*—Time to wait before reconnecting after a session has terminated.

- **Range:** 0 through 4,294,967,295 seconds
- **Default:** 0 (never)

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Configuring PPPoE](#) | 48

## bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

### IN THIS SECTION

- Syntax | 810
- Hierarchy Level | 810
- Description | 810
- Options | 810
- Required Privilege Level | 811
- Release Information | 811

### Syntax

```
bandwidth-limit bps;
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile policer cos-policer-name aggregate],
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile policer cos-policer-name premium]
```

### Description

Define a policer to apply to nonpremium traffic.

### Options

*bps*—Bandwidth limit, in bits per second. Specify either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

- **Range:** 32 Kbps through 32 gigabits per second (Gbps). For IQ2 and IQ2-E interfaces 65,536 bps through 1 Gbps. For 10-Gigabit IQ2 and IQ2-E interfaces 65,536 bps through 10 Gbps.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[burst-size-limit \(Policer for Gigabit Ethernet Interfaces\) | 820](#)

## bert-algorithm

### IN THIS SECTION

- [Syntax | 811](#)
- [Hierarchy Level | 812](#)
- [Description | 812](#)
- [Options | 812](#)
- [Required Privilege Level | 813](#)
- [Release Information | 814](#)

## Syntax

```
bert-algorithm algorithm;
```

## Hierarchy Level

```
[edit interfaces ce1-fpc/pic/port],
[edit interfaces ct1-fpc/pic/port],
[edit interfaces interface-name ds0-options],
[edit interfaces interface-name e1-options],
[edit interfaces interface-name e3-options],
[edit interfaces interface-name t1-options],
[edit interfaces interface-name t3-options]
```

## Description

Configure the pattern to send in the bit stream during a bit error rate test (BERT). Applies to T1, E3, T3, and multichannel DS3 interfaces, the channelized interfaces (DS3, OC12, STM1), and channelized IQ and IQE interfaces (E1, E3 and DS3).

**NOTE:** When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **bert-algorithm** statement must be included at the **[edit interfaces ce1-fpc/pic/port]** or **[edit interfaces ct1-fpc/pic/port]** hierarchy level as appropriate.

## Options

**algorithm**—Pattern to send in the bit stream. There are two categories of test patterns: pseudorandom and repetitive. Both patterns conform to CCITT/ITU O.151, O.152, O.153, and O.161 standards. The algorithm can be one of the following patterns:

- **all-ones-repeating**—Pattern is all ones.
- **all-zeros-repeating**—Pattern is all zeros.
- **alternating-double-ones-zeros**—Pattern is alternating pairs of ones and zeros.
- **alternating-ones-zeros**—Pattern is alternating ones and zeros.
- **pseudo-2e3**—Pattern is  $2^3 - 1$ .
- **pseudo-2e4**—Pattern is  $2^4 - 1$ .
- **pseudo-2e5**—Pattern is  $2^5 - 1$ .
- **pseudo-2e6**—Pattern is  $2^6 - 1$ .

- **pseudo-2e7**—Pattern is  $2^7 - 1$ .
- **pseudo-2e9-o153**—Pattern is  $2^9 - 1$ , as defined in the O153 standard.
- **pseudo-2e10**—Pattern is  $2^{10} - 1$ .
- **pseudo-2e11-o152**—Pattern is  $2^{11} - 1$ , as defined in the O152 standard.
- **pseudo-2e15-o151**—Pattern is  $2^{15} - 1$ , as defined in the O151 standard.
- **pseudo-2e17**—Pattern is  $2^{17} - 1$ .
- **pseudo-2e18**—Pattern is  $2^{18} - 1$ .
- **pseudo-2e20-o151**—Pattern is  $2^{20} - 1$ , as defined in the O151 standard.
- **pseudo-2e20-o153**—Pattern is  $2^{20} - 1$ , as defined in the O153 standard.
- **pseudo-2e21**—Pattern is  $2^{21} - 1$ .
- **pseudo-2e22**—Pattern is  $2^{22} - 1$ .
- **pseudo-2e23-o151**—Pattern is  $2^{23} - 1$ , as defined in the O151 standard.
- **pseudo-2e25**—Pattern is  $2^{25} - 1$ .
- **pseudo-2e28**—Pattern is  $2^{28} - 1$ .
- **pseudo-2e29**—Pattern is  $2^{29} - 1$ .
- **pseudo-2e31**—Pattern is  $2^{31} - 1$ .
- **pseudo-2e32**—Pattern is  $2^{32} - 1$ .
- **repeating-1-in-4**—One bit in four is set to 1; the others are set to 0.
- **repeating-1-in-8**—One bit in eight is set to 1; the others are set to 0.
- **repeating-3-in-24**—Three bits in twenty four are set to 1; the others are set to 0.
- **Default: pseudo-2e3**

### Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

[Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 663](#)

[Configuring E1 BERT Properties](#)

[Configuring E3 BERT Properties](#)

[Configuring T1 BERT Properties](#)

[Configuring T3 BERT Properties](#)

[Examples: Configuring T3 Interfaces](#)

[bert-error-rate | 814](#)

[bert-period | 816](#)

## bert-error-rate

### IN THIS SECTION

- [Syntax | 814](#)
- [Hierarchy Level | 815](#)
- [Description | 815](#)
- [Options | 815](#)
- [Required Privilege Level | 815](#)
- [Release Information | 815](#)

### Syntax

```
bert-error-rate rate;
```

## Hierarchy Level

```
[edit interfaces ce1-fpc/pic/port],
[edit interfaces ct1-fpc/pic/port],
[edit interfaces interface-name ds0-options],
[edit interfaces interface-name e1-options],
[edit interfaces interface-name e3-options],
[edit interfaces interface-name t1-options],
[edit interfaces interface-name t3-options]
```

## Description

Configure the bit error rate to use in a BERT procedure. Applies to E1, E3, T1, or T3 interfaces, and to the channelized interfaces (DS3, OC3, OC12, and STM1).

**NOTE:** When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `bert-error-rate` statement must be included at the `[edit interfaces ce1-fpc/pic/port]` or `[edit interfaces ct1-fpc/pic/port]` hierarchy level as appropriate.

## Options

*rate*—Bit error rate.

- **Range:** 0 through 7, which corresponds to  $10^{-1}$  (1 error per bit) to  $10^{-7}$  (1 error per 10 million bits)
- **Default:** 0

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[bert-algorithm](#) | **811**

[bert-period](#) | **816**

[ds0-options](#)

[e1-options](#)

[e3-options](#)

[t1-options](#)

[t3-options](#)

[Configuring Interface Diagnostics Tools to Test the Physical Layer Connections](#) | **663**

[Configuring E1 BERT Properties](#)

[Configuring E3 BERT Properties](#)

[Configuring T1 BERT Properties](#)

[Configuring T3 BERT Properties](#)

[Examples: Configuring T3 Interfaces](#)

## bert-period

### IN THIS SECTION

- [Syntax](#) | **816**
- [Hierarchy Level](#) | **817**
- [Description](#) | **817**
- [Options](#) | **817**
- [Required Privilege Level](#) | **818**
- [Release Information](#) | **818**

### Syntax

```
bert-period seconds;
```

## Hierarchy Level

```
[edit interfaces ce1-fpc/pic/port],
[edit interfaces ct1-fpc/pic/port],
[edit interfaces interface-name ds0-options],
[edit interfaces interface-name e1-options],
[edit interfaces interface-name e3-options],
[edit interfaces interface-name t1-options],
[edit interfaces interface-name t3-options]
```

## Description

Configure the duration of a BERT test. Applies to E1, E3, T1, and T3 interfaces, and to E1, E3, T1, and T3 partitions on the channelized interfaces (CE1, CT1, DS3, OC3, OC12, OC48, STM1, STM4, and STM16).

E1 and T1 IQ, IQE, and standard interfaces support an extended BERT period range, up to 86,400 seconds (24 hours).

**NOTE:** When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **bert-period** statement must be included at the **[edit interfaces ce1-fpc/pic/port]** or **[edit interfaces ct1-fpc/pic/port]** hierarchy level as appropriate.

## Options

**seconds**—Test duration. Range and default values vary by interface type.

Range:

- • PIC-dependent—Normal BERT period: either 1 through 239 seconds or 1 through 240 seconds
- • PIC-dependent—Extended BERT period: from 1 through 86,400 seconds

Default:

- • Normal BERT period: 10 seconds
- • Extended BERT period (on supported E1 interfaces): 10 seconds
- • Extended BERT period (on supported T1 interfaces): 240 seconds

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 663](#)

[Configuring E1 BERT Properties](#)

[Configuring E3 BERT Properties](#)

[Configuring T1 BERT Properties](#)

[Configuring T3 BERT Properties](#)

[bert-algorithm | 811](#)

[bert-error-rate | 814](#)

## bridge-domain

### IN THIS SECTION

- [Syntax | 819](#)
- [Hierarchy Level | 819](#)
- [Description | 819](#)
- [Options | 819](#)
- [Required Privilege Level | 819](#)
- [Release Information | 819](#)

## Syntax

```
bridge-domain name;
 vlan-id [vlan-identifiers];
}
```

## Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
maintenance-domain-name],
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
maintenance-domain-name virtual-switch virtual-switch-name]
```

## Description

(MX Series routers only) Specify the OAM Ethernet CFM maintenance domain bridge domain.

## Options

***name***—Specify the name of the bridge domain.

***vlan-identifiers***—Specify one or more VLAN identifiers.

## 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 Release 9.4.

## RELATED DOCUMENTATION

*Configuring Maintenance Intermediate Points (MIPs)*

*maintenance-domain*

## burst-size-limit (Policer for Gigabit Ethernet Interfaces)

### IN THIS SECTION

- [Syntax | 820](#)
- [Hierarchy Level | 820](#)
- [Description | 820](#)
- [Options | 820](#)
- [Required Privilege Level | 821](#)
- [Release Information | 821](#)

### Syntax

```
burst-size-limit bytes;
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile policer cos-policer-name aggregate],
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile policer cos-policer-name premium]
```

### Description

Define a policer to apply to nonpremium traffic.

### Options

*bytes*—Burst length.

- **Range:** 1500 through 100,000,000 bytes

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[bandwidth-limit \(Policer for Gigabit Ethernet Interfaces\) | 810](#)

## centralized

### IN THIS SECTION

- [Syntax | 821](#)
- [Hierarchy Level | 822](#)
- [Description | 822](#)
- [Default | 822](#)
- [Required Privilege Level | 822](#)
- [Release Information | 822](#)

## Syntax

```
centralized;
```

## Hierarchy Level

```
[edit protocols lacp ppm]
```

## Description

Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.

This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** statement in the `[edit routing-options ppm]` hierarchy.

**BEST PRACTICE:** We generally recommend that 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.

## Default

Distributed PPM processing is enabled for all packets that use PPM.

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.4.

## RELATED DOCUMENTATION

[inline | 916](#)

[Configuring Distributed Periodic Packet Management on an EX Series Switch \(CLI Procedure\)](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\)](#)

[Configuring Distributed Periodic Packet Management](#)

## ces-psn-channel (tdm-options)

### IN THIS SECTION

- [Syntax | 823](#)
- [Hierarchy Level | 823](#)
- [Description | 823](#)
- [Options | 824](#)
- [Required Privilege Level | 824](#)
- [Release Information | 824](#)

### Syntax

```
ces-psn-channel {
 dmac-address address;
 mode mode;
 vlan-id-1 vlanid;
 vlan-id-2 vlanid;
}
```

### Hierarchy Level

```
[edit interfaces interface-name tdm-options]
```

### Description

Configure the CES (Circuit Emulation Service) PSN (Packet Switched Networks) Channel parameters. Use the parameters to specify the destination MAC address, encapsulation type, and vlan tagging.

## Options

- dmac-address** *address* Destination MAC address to be paired with the smart SFP.
- mode** *mode* Encapsulation mode for the TDM traffic for further network processing. Possible values are: MEF8 and MPLS. Default mode is MEF8 for E1, T1, and DS3 smart SFPs. MPLS encapsulation is not supported for STM1, STM4, and STM16 smart SFPs.
- vlan-id-1** *vlanid* A valid VLAN identifier for single VLAN tagging. Possible values: 0 through 4094.
- vlan-id-2** *vlanid* A valid outer VLAN identifier for dual VLAN tagging. Possible values: 0 through 4094. Dual VLAN tagging is not supported on STM1, STM4, and STM16 smart SFPs.

## 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 Release 19.4.

## RELATED DOCUMENTATION

[show interfaces smart-sfp-statistics | 1719](#)

[show interfaces smart-sfp-defects | 1705](#)

## cfp-to-et

### IN THIS SECTION

- [Syntax | 825](#)
- [Hierarchy Level | 825](#)
- [Description | 825](#)
- [Required Privilege Level | 825](#)

## Syntax

```
cfp-to-et;
```

## Hierarchy Level

```
[edit chassis fpc slot]
```

## Description

Make the interface et-0/1/0 (on the QSFP28 port) available for use. After you configure the **set chassis fpc 0 cfp-to-et** command and commit the configuration, you need to restart the FPC by executing the **restart chassis-control** command. After the FPC comes online, interface et-0/1/0 is created and et-0/2/1(on the CFP2 port) is deleted.

**NOTE:** Before executing this command, plan to handle disruption of services.

## 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 Release 19.2R1-S1.

## RELATED DOCUMENTATION

| [Interface Naming Conventions | 394](#)

## classifier

### IN THIS SECTION

- [Syntax | 826](#)
- [Hierarchy Level | 826](#)
- [Description | 826](#)
- [Required Privilege Level | 827](#)
- [Release Information | 827](#)

### Syntax

```
classifier {
 per-unit-scheduler {
 forwarding-class class-name {
 loss-priority (high | low);
 }
 }
}
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile output-priority-map]
```

### Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the classifier for the output priority map to be applied to outgoing frames on this interface.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[input-priority-map | 920](#)

## client

### IN THIS SECTION

- [Syntax | 827](#)
- [Hierarchy Level | 828](#)
- [Description | 828](#)
- [Required Privilege Level | 828](#)
- [Release Information | 828](#)

## Syntax

```
client;
```

## Hierarchy Level

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

## Description

Configure the router to operate in the PPPoE client mode.

## 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 Release 8.5.

## RELATED DOCUMENTATION

| [Configuring PPPoE](#) | 48

## community-vlans (MX Series)

### IN THIS SECTION

- [Syntax](#) | 829
- [Hierarchy Level](#) | 829
- [Description](#) | 829
- [Options](#) | 829
- [Required Privilege Level](#) | 829
- [Release Information](#) | 830

## Syntax

```
community-vlans [number number-number];
```

## Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains bridge-domain-name ,
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name],
```

## Description

Configure the specified community VLAN to be a secondary VLAN of the specified primary VLAN. A *community* VLAN is used to transport frames among members of a community, which is a subset of users within the VLAN, and to forward frames upstream to the primary VLAN.

**NOTE:** When you specify this configuration statement, the VLAN ID of a logical interface that you associate with a bridge domain that matches with the VLAN ID or list of IDs that you specify using the **community-vlans** state is treated as a community port.

## Options

*number*                    Individual VLAN IDs separated by a space.

## Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 14.2.

## connectivity-fault-management

### IN THIS SECTION

- [Syntax | 830](#)
- [Hierarchy Level | 833](#)
- [Description | 833](#)
- [Options | 833](#)
- [Required Privilege Level | 833](#)
- [Release Information | 833](#)

## Syntax

```
connectivity-fault-management {
 action-profile profile-name {
 action {synchronous-notification
 interface-down;
 log-and-generate-ais {
 interval(1m | 1s);
 levelvalue;
 priority value;
 }
 }
 }
 default-actions {
 interface-down;
 }
 event {
 ais-trigger-condition {
 adjacency-loss;
 all-defects;
 cross-connect-ccm;
```

```

 erroneous-ccm;
 receive-ais;
 }
 adjacency-loss;
 interface-status-tlv (down | lower-layer-down);
 port-status-tlv blocked;
 rdi;
}
}
linktrace {
 age (30m | 10m | 1m | 30s | 10s);
 path-database-size path-database-size;
}
expected-defect {
 rx-enable;
 rx-max-duration seconds;
 tx-enable;
 tx-duration seconds;
}
maintenance-domain domain-name {
 bridge-domain <vlan-id [vlan-ids]>;
 instance routing-instance-name;
 interface interface-name;
 level number;
 name-format (character-string | none | dns | mac+2oct);
 maintenance-association ma-name {
 protect-maintenance-association protect-ma-name;
 remote-maintenance-association remote-ma-name;
 short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-
id);

 continuity-check {
 convey-loss-threshold;
 hold-interval minutes;
 interface-status-tlv;
 interval (10m | 10s | 1m | 1s | 100ms);
 loss-threshold number;
 port-status-tlv;
 }
 mep mep-id {
 auto-discovery;
 direction (up | down);
 interface interface-name (protect | working);
 lowest-priority-defect (all-defects | err-xcon | mac-rem-err-

```

```

xcon | no-defect | rem-err-xcon | xcon);
 priority number;
 remote-mep mep-id {
 action-profile profile-name;
 sla-iterator-profile profile-name {
 data-tlv-size size;
 iteration-count count-value;
 priority priority-value;
 detect-loc;
 }
 }
}
virtual-switch routing-instance-name {
 bridge-domain name <vlan-ids [vlan-ids]>;
}
}
no-aggregate-delegate-processing;
performance-monitoring {
 delegate-server-processing;
 hardware-assisted-timestamping;
 hardware-assisted-keepalives;
 sla-iterator-profiles {
 profile-name {
 avg-fd-twoway-threshold;
 avg-ifdv-twoway-threshold;
 avg-flr-forward-threshold;
 avg-flr-backward-threshold;
 disable;
 calculation-weight {
 delay delay-weight;
 delay-variation delay-variation-weight;
 }
 cycle-time milliseconds;
 iteration-period connections;
 measurement-type (loss | statistical-frame-loss | two-way-delay);
 }
 }
}
}
}

```

## Hierarchy Level

```
[edit protocols oam ethernet]
```

## Description

For Ethernet interfaces on M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M120, M320, MX Series, and T Series routers, specify connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

In Junos OS Release 9.3 and later, this statement is also supported on aggregated Ethernet interfaces.

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

## Options

*Asynchronous-notification*—Generate asynchronous-notification (Laser off for optical physical) option is introduced in Junos OS release 20.4R1 for ACX and MX.

## 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 Release 8.4.

## RELATED DOCUMENTATION

| *IEEE 802.1ag OAM Connectivity Fault Management Overview*

## delay (PPPoE Service Name Tables)

### IN THIS SECTION

- [Syntax | 834](#)
- [Hierarchy Level | 834](#)
- [Description | 834](#)
- [Options | 835](#)
- [Required Privilege Level | 835](#)
- [Release Information | 835](#)

### Syntax

```
delay seconds;
```

### Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-
specifier aci circuit-id-string ari remote-id-string]
```

### Description

Configure the PPPoE underlying interface on the router to wait a specified number of seconds after receiving a PPPoE Active Discovery Initiation (PADI) control packet from a PPPoE client before sending a PPPoE Active Discovery Offer (PADO) packet to indicate that it can service the client request

The router (PPPoE server) does not check whether another server has already sent a PADO packet during the delay period in response to the PPPoE client's PADI packet. It is up to the PPPoE client to determine whether another PPPoE server has responded to its PADI request, or if it must respond to the delayed PADO packet to establish a PPPoE session.

## Options

**seconds**—Number of seconds that the PPPoE underlying interface waits after receiving a PADI packet from a PPPoE client before sending a PADO packet in response.

- **Range:** 1 through 120 seconds

## 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 Release 10.0.

Support at [edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*] hierarchy level introduced in Junos OS Release 10.2.

## RELATED DOCUMENTATION

| *Configuring PPPoE Service Name Tables*

## destination (IPCP)

### IN THIS SECTION

- [Syntax | 836](#)
- [Hierarchy Level | 836](#)
- [Description | 836](#)
- [Options | 836](#)
- [Required Privilege Level | 836](#)
- [Release Information | 836](#)

## Syntax

```
destination address destination-profile profile-name;
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet unnumbered-
address interface-name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family inet unnumbered-address interface-name]
```

## Description

For unnumbered interfaces with PPP encapsulation, specify the IP address of the remote interface.

## Options

***address***—IP address of the remote interface.

The remaining statement is explained separately. See [CLI Explorer](#).

## Required Privilege Level

**interface**—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring IPCP Options for Interfaces with PPP Encapsulation](#)

[address](#) | [791](#)

[negotiate-address](#) | [1025](#)

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

## device-count

### IN THIS SECTION

- [Syntax | 837](#)
- [Hierarchy Level | 837](#)
- [Description | 837](#)
- [Options | 838](#)
- [Required Privilege Level | 838](#)
- [Release Information | 838](#)

### Syntax

```
device-count number;
```

### Hierarchy Level

```
[edit chassis aggregated-devices ethernet]
[edit chassis aggregated-devices sonet]
```

### Description

Configure the number of aggregated logical devices available to the router.

Starting in Junos release 14.2, for MX series routers, aggregated Ethernet interfaces created under a logical system can be individually named. Prior to 14.2, ae interfaces were named automatically (AE1, AE2) etc. upon setting the device count. This change allows administrators to use custom naming schemes. System resources are only allocated for named ae interfaces, regardless of how many were declared in the device count. (In Junos 14.2 and earlier, ae naming occurred automatically up to the number specified for device count, and system resources were allocated whether a given ae interface was used or not.)

## Options

*number*—Set the number of aggregated logical devices that will be available for configuration.

**NOTE:** Starting with Junos OS Release 13.2, a maximum of 64 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces. In releases before Junos OS Release 13.2, a maximum of 16 aggregated interfaces are supported for link aggregation of SONET/SDH interfaces.

For Junos OS Evolved, you can specify up to 512 aggregated Ethernet devices.

- **Range:** 1 - 496. The upper limit for this value is system specific.
- **Range:** 1 - 512 for Junos OS Evolved.

## Required Privilege Level

*interface*—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement functionality updated in Junos OS Release 14.2, as described below.

## RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices](#)

[Configuring Aggregated SONET/SDH Interfaces](#)

## drop (PPPoE Service Name Tables)

### IN THIS SECTION

● [Syntax | 839](#)

- [Hierarchy Level | 839](#)
- [Description | 839](#)
- [Required Privilege Level | 839](#)
- [Release Information | 839](#)

## Syntax

```
drop;
```

## Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-
specifier aci circuit-id-string ari remote-id-string]
```

## Description

Direct the router to drop (ignore) a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client that contains the specified service name tag or agent circuit identifier/agent remote identifier (ACI/ARI) information. This action effectively denies the client's request to provide the specified service, or to accept requests from the subscriber or subscribers represented by the ACI/ARI information.

## 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 Release 10.0.

Support at [edit protocols pppoe service-name-tables *table-name* service *service-name* agent-specifier aci *circuit-id-string* ari *remote-id-string*] hierarchy level introduced in Junos OS Release 10.2.

## RELATED DOCUMENTATION

| [Configuring PPPoE Service Name Tables](#)

## dynamic-profile (PPPoE Service Name Tables)

### IN THIS SECTION

- [Syntax | 840](#)
- [Hierarchy Level | 840](#)
- [Description | 841](#)
- [Options | 841](#)
- [Required Privilege Level | 841](#)
- [Release Information | 841](#)

### Syntax

```
dynamic-profile profile-name;
```

### Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-specifier
aci circuit-id-string ari remote-id-string]
```

## Description

Specify a dynamic profile to instantiate a dynamic PPPoE interface. You can associate a dynamic profile with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The dynamic profile associated with a service entry in a PPPoE service name table overrides the dynamic profile associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **dynamic-profile** statement at the `[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci circuit-id-string ari remote-id-string]` hierarchy level, you cannot also include the **static-interface** statement at this level. The **dynamic-profile** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

## Options

***profile-name***—Name of the dynamic profile that the router uses to instantiate a dynamic PPPoE interface.

## 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 Release 10.2.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

*Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

## egress-policer-overhead

### IN THIS SECTION

- [Syntax | 842](#)
- [Hierarchy Level | 842](#)
- [Description | 842](#)
- [Options | 843](#)
- [Required Privilege Level | 843](#)
- [Release Information | 843](#)

### Syntax

```
egress-policer-overhead bytes;
```

### Hierarchy Level

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

### Description

Add the specified number of bytes to the actual length of an Ethernet frame when determining the actions of Layer 2 policers, MAC policers, or queue rate limits applied to output traffic on the line card. You can configure egress policer overhead to account for egress *shaping* overhead bytes added to output traffic on the line card.

On M Series and T Series routers, this statement is supported on Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs and Enhanced IQ2 (IQ2E) PICs. On MX Series routers, this statement is supported for interfaces configured on Dense Port Concentrators (DPCs).

**NOTE:** This statement is not supported on Modular Interface Cards (MICs) or Modular Port Concentrators (MPCs) in MX Series routers.

## Options

*bytes*—Number of bytes added to a packet exiting an interface.

- **Range:** 0–255 bytes
- **Default:** 0

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 11.1.

## RELATED DOCUMENTATION

*egress-shaping-overhead*

*Policer Overhead to Account for Rate Shaping Overview*

*Example: Configuring Policer Overhead to Account for Rate Shaping*

*Configuring a Policer Overhead*

*CoS on Enhanced IQ2 PICs Overview*

## encapsulation (Logical Interface)

### IN THIS SECTION

- [Syntax | 844](#)
- [Hierarchy Level | 844](#)
- [Description | 844](#)
- [Options | 844](#)
- [Required Privilege Level | 847](#)

## Syntax

```
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-tcc-vc-mux | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | ethernet-ccc | ethernet-vpls | ethernet-vpls-fr | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-ppp | frame-relay-tcc | gre-fragmentation | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls | vxlan);
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
[edit interfaces rlsq number unit logical-unit-number]
[edit protocols evpn]
```

## Description

Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. See [Configuring VLAN TCC Encapsulation](#) for more details. Non-ethernet media types, SONET and ATM interfaces are only supported. It is expected that the user will have the member links of aggregated ethernet with supported hardware for configuring VLAN TCC encapsulation and no commit check is performed externally for the aggregated ethernet (AE) interfaces.

## Options

**atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.

**atm-ccc-vc-mux**—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**atm-cisco-nlpid**—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

**atm-mlppp-llc**—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

**atm-nlpid**—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

**atm-ppp-llc**—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.

**atm-ppp-vc-mux**—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.

**atm-snap**—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

**atm-tcc-snap**—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

**atm-tcc-vc-mux**—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

**atm-vc-mux**—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

**ether-over-atm-llc**—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

**ether-vpls-over-atm-llc**—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

**ether-vpls-over-fr**—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.

**NOTE:** The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.

**ether-vpls-over-ppp**—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

**ethernet**—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

**ethernet-ccc**—Use Ethernet CCC encapsulation on Ethernet interfaces.

**ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

**NOTE:** The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

**ethernet-vpls-fr**—Use in a VPLS setup when a CE device is connected to a PE router over a time-division multiplexing (TDM) link. This encapsulation type enables the PE router to terminate the outer layer 2 Frame Relay connection, use the *802.1p* bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

**frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

**frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

**frame-relay-ppp**—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

**frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

**gre-fragmentation**—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF)

bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

**multilink-frame-relay-end-to-end**—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

**multilink-ppp**—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

**ppp-over-ether**—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

**ppp-over-ether-over-atm-llc**—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

**vlan-bridge**—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE *802.1Q* tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

**vlan-ccc**—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**vlan-vci-ccc**—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

**vlan-tcc**—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

**vlan-vpls**—Use Ethernet VLAN encapsulation on VPLS circuits.

**vxlan**—Use VXLAN data plane encapsulation for EVPN.

## Required Privilege Level

**interface**—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## Release History Table

| Release | Description                                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20.1R1  | Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. |

## RELATED DOCUMENTATION

- [Configuring Layer 2 Switching Cross-Connects Using CCC](#)
- [Configuring the Encapsulation for Layer 2 Switching TCCs](#)
- [Configuring Interface Encapsulation on Logical Interfaces](#)
- [Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects](#)
- [Circuit and Translational Cross-Connects Overview](#)
- [Identifying the Access Concentrator](#)
- [Configuring ATM Interface Encapsulation](#)
- [Configuring VLAN and Extended VLAN Encapsulation](#)
- [Configuring ATM-to-Ethernet Interworking](#)
- [Configuring Interface Encapsulation on PTX Series Packet Transport Routers](#)
- [\*Configuring CCC Encapsulation for Layer 2 VPNs\*](#)
- [\*Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits\*](#)
- [Configuring ATM for Subscriber Access](#)
- [Understanding CoS on ATM IMA Pseudowire Interfaces Overview](#)
- [Configuring Policing on an ATM IMA Pseudowire](#)

## encapsulation

### IN THIS SECTION

- [Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series | 849](#)
- [Syntax for Physical Interfaces: SRX Series | 849](#)
- [Syntax for Logical Interfaces: SRX Series | 849](#)
- [Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series | 850](#)

- [Logical Interfaces | 850](#)
- [Description | 850](#)
- [Default | 850](#)
- [Physical Interface Options and Logical Interface Options | 850](#)
- [Required Privilege Level | 855](#)
- [Release Information | 855](#)

## Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

```
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc |
cisco-hdlc-tcc | ethernet-bridge | ethernet-ccc | ethernet-over-atm | ethernet-
tcc | ethernet-vpls | ethernet-vpls-fr | ether-vpls-over-atm-llc | ethernet-vpls-
ppp | extended-frame-relay-ccc | extended-frame-relay-ether-type-tcc | extended-
frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc |
extended-vlan-vpls | flexible-ethernet-services | flexible-frame-relay | frame-
relay | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc |
frame-relay-port-ccc | frame-relay-tcc | generic-services | multilink-frame-
relay-uni-nni | ppp | ppp-ccc | ppp-tcc | vlan-ccc | vlan-vci-ccc | vlan-vpls);
```

## Syntax for Physical Interfaces: SRX Series

```
encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc |
ethernet-vpls | extended-frame-relay-ccc | extended-frame-relay-tcc | extended-
vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls |
flexible-ethernet-services | frame-relay-port-ccc | vlan-ccc | vlan-vpls);
```

## Syntax for Logical Interfaces: SRX Series

```
encapsulation (dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-
bridge | vlan-ccc | vlan-tcc | vlan-vpls);
```

## Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

```
[edit interfaces interface-name],
[edit interfaces rlsq number:number]
```

## Logical Interfaces

```
[edit interfaces interface-name unit logical-
unit-number]
```

## Description

For M Series, MX Series, QFX Series, T Series, PTX Series, specify the physical link-layer encapsulation type.

For SRX Series, specify logical link layer encapsulation.

**NOTE:** Not all encapsulation types are supported on the switches. See the switch CLI.

## Default

**ppp**—Use serial PPP encapsulation.

## Physical Interface Options and Logical Interface Options

For physical interfaces:

**NOTE:** Frame Relay, ATM, PPP, SONET, and SATSOP options are not supported on EX Series switches.

- **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.
- **atm-pvc**—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label

Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).

- **cisco-hdlc**—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:
  - CCC version (**cisco-hdlc-ccc**)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
  - TCC version (**cisco-hdlc-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **cisco-hdlc-ccc**—Use Cisco-compatible HDLC framing on CCC circuits.
- **cisco-hdlc-tcc**—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.
- **ethernet-bridge**—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.
- **ethernet-over-atm**—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.
- **ethernet-tcc**—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.
- **ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.
- **ethernet-vpls-fr**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the *802.1p* bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

- **ethernet-vpls-ppp**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.
- **ether-vpls-over-atm-llc**—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.
- **extended-frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. When you use this encapsulation type, you can configure the **ccc** family only.
- **extended-frame-relay-ether-type-tcc**—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.
- **extended-frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.
- **extended-vlan-bridge**—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE *802.1Q* VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.
- **extended-vlan-ccc**—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.
- **extended-vlan-tcc**—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.
- **extended-vlan-vpls**—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

**NOTE:** The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

- **flexible-ethernet-services**—For Gigabit Ethernet IQ interfaces and 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), and for Gigabit Ethernet interfaces, use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.
- **flexible-frame-relay**—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.
- **frame-relay**—Use Frame Relay encapsulation is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation.
- **frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation is same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- **frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload.

**NOTE:** When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

- **frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. This encapsulation is Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC.
- **frame-relay-port-ccc**—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers

can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. When you use this encapsulation type, you can configure the **ccc** family only.

- **frame-relay-tcc**—This encapsulation is similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **generic-services**—Use generic services encapsulation for services with a hierarchical scheduler.
- **multilink-frame-relay-uni-nni**—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.
- 
- **ppp**—Use serial PPP encapsulation. This encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.
- **ppp-ccc**—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.
- **ppp-tcc**—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.
- **vlan-ccc**—Use Ethernet VLAN encapsulation on CCC circuits. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.
- **vlan-vci-ccc**—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.
- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

**NOTE:**

- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.

- Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure **family inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.

For logical interfaces:

- **frame-relay**—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint.
- **multilink-frame-relay-uni-nni**—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation.
- **ppp**—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface.
- **ppp-over-ether**—This encapsulation is used for underlying interfaces of pp0 interfaces.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5.

## RELATED DOCUMENTATION

[Understanding Physical Encapsulation on an Interface](#)

[Configuring Interface Encapsulation on Physical Interfaces](#)

*Configuring CCC Encapsulation for Layer 2 VPNs*

[Configuring Layer 2 Switching Cross-Connects Using CCC](#)

*Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*

[Configuring ATM Interface Encapsulation](#)

[Configuring ATM-to-Ethernet Interworking](#)

[Configuring VLAN and Extended VLAN Encapsulation](#)

[Configuring VLAN and Extended VLAN Encapsulation](#)

[Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces](#)

*Configuring Interfaces for Layer 2 Circuits*

[Configuring Interface Encapsulation on PTX Series Packet Transport Routers](#)

[Configuring MPLS LSP Tunnel Cross-Connects Using CCC](#)

[Configuring TCC](#)

*Configuring VPLS Interface Encapsulation*

*Configuring Interfaces for VPLS Routing*

[Defining the Encapsulation for Switching Cross-Connects](#)

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\)](#)

## ether-options

### IN THIS SECTION

- [Junos OS Syntax | 856](#)
- [Junos OS Evolved Syntax | 857](#)
- [Hierarchy Level | 858](#)
- [Description | 858](#)
- [Default | 864](#)
- [Options | 864](#)
- [Required Privilege Level | 865](#)
- [Release Information | 865](#)

## Junos OS Syntax

```
ether-options {
 802.3ad {
 aex;
 (backup | primary);
 }
 lacp {
 force-up;
 (primary | backup);
 }
}
```

```

 port-priority
 }
}
asynchronous-notification;
(auto-negotiation| no-auto-negotiation);

configured-flow-control {
 rx-buffers (on | off);
 tx-buffers (on | off);
}
ethernet-switch-profile {
 ethernet-policer-profile
 (mac-learn-enable | no-mac-learn-enable);
 recovery-timeout time-in-seconds;
 storm-control storm-control-profile;
 tag-protocol-id;
}
(flow-control | no-flow-control);
ieee-802-3az-eee;
ignore-l3-incompletes;
link-mode (automatic | full-duplex | half-duplex);
(loopback | no-loopback);
mdi-mode (auto | force | mdi | mdix);
mpls {
 pop-all-labels <required-depth (1 | 2 | all)>;
}
no-auto-mdix;
redundant-parent (Interfaces) parent;
source-address-filter name;
(source-filtering| no-source-filtering);
speed {
 (auto-negotiation <auto-negotiate-10-100> | ethernet-100m | ethernet-10g
| ethernet-10m | ethernet-1g);
}
}

```

## Junos OS Evolved Syntax

```

ether-options {
 802.3ad {
 aex;
 }
}

```

```

 (backup | primary);
 lacp {
 force-up;
 (primary | backup);
 port-priority
 }
}
asynchronous-notification;
(auto-negotiation| no-auto-negotiation);

ethernet-switch-profile {
 ethernet-policer-profile
 (mac-learn-enable | no-mac-learn-enable);
 recovery-timeout time-in-seconds;
 storm-control storm-control-profile;
 tag-protocol-id;
}
fec (gigether)
(flow-control | no-flow-control);
ignore-l3-incompletes;
(loopback | no-loopback);
loopback-remote;
mpls {
 pop-all-labels <required-depth (1 | 2 | all)>;
}
source-address-filter name;
(source-filtering | no-source-filtering);
}

```

## Hierarchy Level

```

[edit interfaces interface-name]
[edit interfaces interface-range range]

```

## Description

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

In Junos OS Evolved, when you configure **set interfaces *interface* ether-options 802.3ad *ae name*** at the same time as you apply a second configuration to the same interface at the **[edit interfaces *interface*]**

hierarchy, the second configuration will not take effect until the interface joins the aggregated Ethernet interface *ae name*.

**NOTE:** The **ether-options** statement is not supported for subscriber management on aggregated Ethernet member link interfaces. You must configure **gether-options** instead.

Table 128 on page 859 shows the supported and unsupported platforms.

**Table 128: Supported Platform Information**

| Supported Platforms for ggether-options                                                                                                                                                                                                                                               | Supported Platforms for ether-options                                                                        | Notes                                                                            |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| ACX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• ACX500</li> <li>• ACX710</li> <li>• ACX1000</li> <li>• ACX1100</li> <li>• ACX2100</li> <li>• ACX2200</li> <li>• ACX4000</li> <li>• ACX5400</li> <li>• ACX5448</li> <li>• ACX5048</li> <li>• ACX5096</li> </ul> | ACX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• ACX5048</li> <li>• ACX5096</li> </ul> | To configure storm control on the ACX Series routers, use <b>ether-options</b> . |

Table 128: Supported Platform Information (*Continued*)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                                                                                                        | Supported Platforms for <b>ether-options</b>                                                                                                                                                                                                                                           | Notes                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• EX2300</li> <li>• EX2300 Multigigabit</li> <li>• EX2300-C</li> <li>• EX3400</li> <li>• EX4300</li> <li>• EX4300 Multigigabit</li> <li>• EX4600</li> <li>• EX4650</li> <li>• EX9200</li> <li>• EX9250</li> </ul> | EX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• EX2300</li> <li>• EX2300 Multigigabit</li> <li>• EX2300-C</li> <li>• EX3400</li> <li>• EX4300</li> <li>• EX4300 Multigigabit</li> <li>• EX4600</li> <li>• EX4650</li> <li>• EX9200</li> <li>• EX9250</li> </ul> | EX Series Switches (Junos OS)<br><br>EX Series switches support both <b>ether-options</b> and <b>gigether-options</b> .<br><br>To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use <b>ether-options</b> .<br><br>To configure Forward Error Correction (FEC), use <b>gigether-options</b> . |

Table 128: Supported Platform Information (*Continued*)

| Supported Platforms for gigheter-options                                                                                                                                                                                                                                                                                                            | Supported Platforms for ether-options | Notes |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|-------|
| MX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• MX5</li> <li>• MX10</li> <li>• MX40</li> <li>• MX80</li> <li>• MX104</li> <li>• MX150</li> <li>• MX204</li> <li>• MX240</li> <li>• MX480</li> <li>• MX960</li> <li>• MX20008</li> <li>• MX2010</li> <li>• MX2020</li> <li>• MX10003</li> <li>• MX10008 and MX10016</li> </ul> | Not Supported                         | None  |

Table 128: Supported Platform Information (Continued)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                          | Supported Platforms for <b>ether-options</b>                                                                                                                       | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PTX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• PTX1000</li> <li>• PTX3000</li> <li>• PTX5000</li> <li>• PTX10001</li> <li>• PTX10002</li> <li>• PTX10008 and PTX10016</li> </ul> | PTX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• PTX1000</li> <li>• PTX10001</li> <li>• PTX10002</li> <li>• PTX10008 and PTX10016</li> </ul> | PTX Series Routers (Junos OS) <p>In Junos OS Release 17.3R3S7, PTX1000 Series routers support both <b>ether-options</b> and <b>gigether-options</b>.</p> <p>In Junos OS Releases 17.3R1, 17.4R1, and 17.4R2, PTX10000 Series routers support both <b>ether-options</b> and <b>gigether-options</b>.</p> <p>In Junos OS Release 18.1R1 and later, PTX Series routers support only <b>gigether-options</b> for your configuration and do not support <b>ether-options</b>. Before Junos OS Release 18.1R1, PTX Series support <b>ether-options</b>. See <a href="#">TSB17864</a> for additional information.</p> |
| PTX Series routers (Junos OS Evolved) <ul style="list-style-type: none"> <li>• PTX10003</li> <li>• PTX10008 and PTX10016</li> </ul>                                                                      | PTX Series routers (Junos OS Evolved) <ul style="list-style-type: none"> <li>• PTX10003</li> <li>• PTX10008 and PTX10016</li> </ul>                                | PTX Series routers (Junos OS Evolved) <p>Starting in Junos OS Evolved Release 20.1R1, PTX Series routers support <b>ether-options</b> only.</p> <p>Before Junos OS Evolved Release 20.1R1, PTX Series routers supported <b>gigether-options</b> only.</p>                                                                                                                                                                                                                                                                                                                                                      |

Table 128: Supported Platform Information (*Continued*)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Supported Platforms for <b>ether-options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Notes                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| QFX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• QFX5100 (48S)</li> <li>• QFX5100 (48T)</li> <li>• QFX5100 (24Q)</li> <li>• QFX5100 (96S)</li> <li>• QFX5110 (48S)</li> <li>• QFX5110 (32Q)</li> <li>• QFX5120 (48Y)</li> <li>• QFX5120 (48T)</li> <li>• QFX5120 (48ym)</li> <li>• QFX5120 (32C)</li> <li>• QFX5200 (48Y)</li> <li>• QFX5200 (32C)</li> <li>• QFX5210 (64C)</li> <li>• QFX5210 (64C-S)</li> <li>• QFX5220 (32CD)</li> <li>• QFX5220 (128C)</li> <li>• QFX10002</li> <li>• QFX10008 and QFX10016</li> </ul> | QFX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• QFX5100 (48S)</li> <li>• QFX5100 (48T)</li> <li>• QFX5100 (24Q)</li> <li>• QFX5100 (96S)</li> <li>• QFX5110 (48S)</li> <li>• QFX5110 (32Q)</li> <li>• QFX5120 (48Y)</li> <li>• QFX5120 (48T)</li> <li>• QFX5120 (48ym)</li> <li>• QFX5120 (32C)</li> <li>• QFX5200 (48Y)</li> <li>• QFX5200 (32C)</li> <li>• QFX5210 (64C)</li> <li>• QFX5210 (64C-S)</li> <li>• QFX5220 (32CD)</li> <li>• QFX5220 (128C)</li> <li>• QFX10002</li> <li>• QFX10008 and QFX10016</li> </ul> | QFX Series Switches (Junos OS) <p>QFX Series switches support both <b>ether-options</b> and <b>gigether-options</b>.</p> <p>To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use <b>ether-options</b>.</p> <p>To configure Forward Error Correction (FEC), use <b>gigether-options</b>.</p> |

Table 128: Supported Platform Information (*Continued*)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                               | Supported Platforms for <b>ether-options</b>                                                                                                                                                                  | Notes                                                                                                                                                                                                          |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Not Supported                                                                                                                                                                                                 | QFX Series Switches (Junos OS Evolved) <ul style="list-style-type: none"> <li>• QFX5200-32C-L</li> <li>• QFX5220-32CD</li> <li>• QFX5220-128C</li> </ul>                                                      | To configure auto-negotiation on management interfaces ( <b>re0:mgmt-0</b> ), use <b>ether-options</b> .                                                                                                       |
| SRX Series (Junos OS) <ul style="list-style-type: none"> <li>• SRX300</li> <li>• SRX550</li> <li>• SRX1500</li> <li>• SRX4100 and SRX4200</li> <li>• SRX4600</li> <li>• SRX5400</li> <li>• SRX5600</li> </ul> | SRX Series (Junos OS) <ul style="list-style-type: none"> <li>• SRX300</li> <li>• SRX550</li> <li>• SRX1500</li> <li>• SRX4100 and SRX4200</li> <li>• SRX4600</li> <li>• SRX5400</li> <li>• SRX5600</li> </ul> | SRX Series (Junos OS) <p>To configure gigabit-Ethernet interfaces (ge-), use <b>gigether-options</b>. To configure ethernet interfaces (et-) and fast ethernet interfaces (fe-), use <b>ether-options</b>.</p> |

## Default

Enabled.

## Options

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

**loopback-remote** Starting in Junos OS Evolved Release 20.1R1, enable remote loopback.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.0.

**autostate-exclude** option introduced in Junos OS Release 14.1x53-D40 for QFX5100 switches only.

**fec** and **loopback-remote** options introduced in Junos OS Evolved Release 20.1R1.

## RELATED DOCUMENTATION

*Gigabit Ethernet Interface*

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

[Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support](#)

[gigether-options](#) | 896

## ethernet (Chassis)

### IN THIS SECTION

- [Syntax](#) | 866
- [Hierarchy Level](#) | 866
- [Description](#) | 866
- [Required Privilege Level](#) | 866
- [Release Information](#) | 866

## Syntax

```
ethernet {
 device-count number;
 lacp {
 link-protection {
 non-revertive;
 }
 system-priority;
 }
}
```

## Hierarchy Level

```
[edit chassis aggregated-devices]
```

## Description

Configure properties for Ethernet aggregated devices on the router.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## ethernet-policer-profile

### IN THIS SECTION

- [Syntax | 867](#)
- [Hierarchy Level | 868](#)
- [Description | 868](#)
- [Required Privilege Level | 868](#)
- [Release Information | 868](#)

### Syntax

```
ethernet-policer-profile {
 input-priority-map {
 ieee802.1p premium [values];
 }
 output-priority-map {
 classifier {
 premium {
 forwarding-class class-name {
 loss-priority (high | low);
 }
 }
 }
 }
 policer cos-policer-name {
 aggregate {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 premium {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 }
}
```

## Hierarchy Level

```
[edit interfaces interface-name gigheter-options ethernet-switch-profile],
[edit interfaces interface-name aggregated-ether-options ethernet-switch-profile]
```

## Description

**NOTE:** On QFX Series standalone switches, this statement hierarchy is only supported on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, configure a class of service (CoS)-based policer. Policing applies to the inner VLAN identifiers, not to the outer tag. For Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the **premium** policer is not supported.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Configuring Gigabit Ethernet Policers](#) | 271

## evcs

### IN THIS SECTION

- [Syntax | 869](#)
- [Hierarchy Level | 869](#)
- [Description | 869](#)
- [Options | 869](#)
- [Required Privilege Level | 870](#)
- [Release Information | 870](#)

### Syntax

```
evcs evc-id {
 evc-protocol cfm;
 remote-uni-count count;
 multipoint-to-multipoint;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet]
```

### Description

On MX Series routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet virtual connection.

### Options

**remote-uni-count** *count*—(Optional) Specify the number of remote UNIs in the EVC configuration, the default is 1.

**multipoint-to-multipoint**—(Optional) Specify multiple points in the EVC configuration, the default is point-to-point if **remote-uni-count** is 1.

Remaining options are explained separately.

### 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 Release 9.5.

### RELATED DOCUMENTATION

*Configuring Ethernet Local Management Interface*

[Imi \(Ethernet OAM\) | 973](#)

## family

### IN THIS SECTION

- [Syntax | 870](#)
- [Hierarchy Level | 871](#)
- [Description | 871](#)
- [Options | 871](#)
- [Required Privilege Level | 871](#)
- [Release Information | 872](#)

### Syntax

```
family {
 inet {
 layer-3;
 layer-4;
```

```

 symmetric-hash {
 complement;
 }
 }
 multiservice {
 source-mac;
 destination-mac;
 payload {
 ip {
 layer-3;
 layer-4;
 }
 }
 symmetric-hash {
 complement;
 }
 }
}

```

## Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key]
```

## Description

(MX Series 5G Universal Routing Platforms only) Configure data used in a hash key for a specific protocol family when configuring PIC-level symmetrical load balancing on an 802.3ad Link Aggregation Group.

## Options

- inet**                Configure data used in a hash key for the **inet** protocol family.
- multiservice**        Configure data used in a hash key for the **multiservice** protocol family.

## 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 Release 9.6.

## RELATED DOCUMENTATION

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

## family

### IN THIS SECTION

- [Syntax | 872](#)
- [Hierarchy Level | 875](#)
- [Description | 876](#)
- [Options | 876](#)
- [Required Privilege Level | 877](#)
- [Release Information | 878](#)

## Syntax

```
family family {
 accounting {
 destination-class-usage;
 source-class-usage {
 (input | output | input output);
 }
 }
 access-concentrator name;
 address address {
 ... the address subhierarchy appears after the main [edit interfaces
```

```

interface-name unit logical-unit-number family family-name] hierarchy ...
}
bundle interface-name;
core-facing;
demux-destination {
 destination-prefix;
}
demux-source {
 source-prefix;
}
direct-connect;
duplicate-protection;
dynamic-profile profile-name;
filter {
 group filter-group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
}
interface-mode (access | trunk);
ipsec-sa sa-name;
keep-address-and-control;
mac-validate (loose | strict);
max-sessions number;
max-sessions-vs-a-ignore;
mtu bytes;
multicast-only;
nd6-stale-time seconds;
negotiate-address;
no-neighbor-learn;
no-redirects;
policer {
 arp policer-template-name;
 input policer-template-name;
 output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);

```

```

rpf-check {
 fail-filter filter-name
 mode loose;
}
sampling {
 input;
 output;
}
service {
 input {
 post-service-filter filter-name;
 service-set service-set-name <service-filter filter-name>;
 }
 output {
 service-set service-set-name <service-filter filter-name>;
 }
}
service-name-table table-name;
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
maximum-seconds> <filter [aci]>;
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address destination-
profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
address address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 destination address;
 destination-profile name;
 eui-64;
 primary-only;
 multipoint-destination address dlci dlci-identifier;
 multipoint-destination address {
 epd-threshold cells;
 inverse-arp;
 oam-liveness {
 up-count cells;
 down-count cells;
 }
 oam-period (disable | seconds);
 }
}

```

```

 shaping {
 (cbr rate | rtvbr burst length peak rate sustained rate | vbr
burst length peak rate sustained rate);
 queue-length number;
 }
 vci vpi-identifier.vci-identifier;
 }
 preferred;
 primary;
 vrrp-group group-id {
 (accept-data | no-accept-data);
 advertise-interval seconds;
 authentication-key key;
 authentication-type authentication;
 fast-interval milliseconds;
 (preempt | no-preempt) {
 hold-time seconds;
 }
 priority number;
 track {
 interface interface-name {
 bandwidth-threshold bits-per-second priority-cost priority;
 priority-cost priority;
 }
 priority-hold-time seconds;
 route prefix routing-instance instance-name priority-cost
priority;
 }
 }
 virtual-address [addresses];
}
virtual-link-local-address ipv6-address;
}
}

```

## Hierarchy Level

```

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

```

## Description

Configure protocol family information for the logical interface.

Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. See [Configuring VLAN TCC Encapsulation](#) for more details.

**NOTE:** Not all subordinate statements are available to every protocol family.

## Options

*family*—Protocol family:

- **any**—Protocol-independent family used for Layer 2 packet filtering

**NOTE:** This option is not supported on T4000 Type 5 FPCs.

- **bridge**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation. You can optionally configure this protocol family for the logical interface on which you configure VPLS.
- **ethernet-switching**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- **ccc**—Circuit cross-connect protocol suite. You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the **ccc** family only.
- **inet**—Internet Protocol version 4 suite. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—Internet Protocol version 6 suite. You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP).

- **iso**—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite. You must configure this protocol family for the logical interface to support IS-IS traffic.
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16. You must configure this protocol or **mlfr-end-to-end** for the logical interface to support link services and voice services bundling.
- **multilink-ppp**—Multilink Point-to-Point Protocol. You must configure this protocol (or **mlfr-end-to-end**) for the logical interface to support multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- **pppoe**—Point-to-Point Protocol over Ethernet
- **tcc**—Translational cross-connect protocol suite. You can configure this protocol family for the logical interface of TCC physical interfaces.
- **tnp**—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only, as discussed in [Understanding Internal Ethernet Interfaces](#).
- **vpls**—(M Series and T Series routers only) Virtual private LAN service. You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the **family vpls** statement is assumed by default.

MX Series routers support dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration.

For more information about VPLS, see the [Junos OS VPNs Library for Routing Devices](#).

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

## Required Privilege Level

**interface**—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Option **max-sessions-vs-a-ignore** introduced in Junos OS Release 11.4.

### Release History Table

| Release | Description                                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20.1R1  | Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. |

### RELATED DOCUMENTATION

| [Configuring the Protocol Family](#)

## fastether-options

### IN THIS SECTION

- [Syntax | 878](#)
- [Hierarchy Level | 879](#)
- [Description | 879](#)
- [Required Privilege Level | 879](#)
- [Release Information | 879](#)

## Syntax

```
fastether-options {
 802.3ad {
 aex (primary | backup);
 lacp {
 port-priority;
 }
 }
}
```

```

}
(flow-control | no-flow-control);
ignore-l3-incompletes;
ingress-rate-limit rate;
(loopback | no-loopback);
mpls {
 pop-all-labels {
 required-depth number;
 }
}
source-address-filter {
 mac-address;
}
(source-filtering | no-source-filtering);
}

```

## Hierarchy Level

```
[edit interfaces interface-name]
```

## Description

Configure Fast Ethernet-specific interface properties.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Ethernet Interfaces Overview](#) | 2

## flow-control

### IN THIS SECTION

- [Syntax | 880](#)
- [Hierarchy Level | 880](#)
- [Description | 880](#)
- [Default | 881](#)
- [Required Privilege Level | 881](#)
- [Release Information | 881](#)

### Syntax

```
(flow-control | no-flow-control);
```

### Hierarchy Level

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options],
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options],
[edit interfaces interface-name multiservice-options],
[edit interfaces interface-range name aggregated-ether-options],
[edit interfaces interface-range name ether-options]
```

### Description

For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.

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

## Default

Flow control is enabled.

**NOTE:** Flow control is enabled by default only on physical interfaces and it is disabled by default on aggregated Ethernet interfaces.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

## RELATED DOCUMENTATION

[Configuring Flow Control | 18](#)

[Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

*Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support*

## fnp

### IN THIS SECTION

- [Syntax | 882](#)
- [Hierarchy Level | 882](#)
- [Description | 882](#)
- [Options | 882](#)
- [Required Privilege Level | 883](#)
- [Release Information | 883](#)

### Syntax

```
fnp {
 interval <100ms | 1s | 10s | 1m | 10m>;
 loss-threshold number
 interface interface name {
 domain-id domain-id
 }
}
```

### Hierarchy Level

```
[edit protocols oam ethernet]
```

### Description

On routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet failure notification protocol.

### Options

**interval *number***—Specifies the time between the transmission of FNP messages.

**loss-threshold *number***—FNP messages that can be lost before the FNP message is considered aged out and flushed.

**interface *interface-name***—Name of the Ethernet interface.

**domain-id *number***—Domain ID of the access network.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Command introduced in Junos OS Release 11.4.

## RELATED DOCUMENTATION

*Ethernet Failure Notification Protocol Overview*

[Configuring the Failure Notification Protocol](#)

## force-up

### IN THIS SECTION

- [Syntax | 884](#)
- [Hierarchy Level | 884](#)
- [Description | 884](#)
- [Required Privilege Level | 884](#)
- [Release Information | 884](#)

## Syntax

```
force-up;
```

## Hierarchy Level

```
[edit interfaces interface-name aggregated-ethernet-options lacp]
```

## Description

Configure the peer interface (in MC-LAG) to remain up even with limited LACP capability.

## 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 Release 14.2.

**NOTE:** For EX9200 switches, you must configure **force-up** on physical interfaces of both MC-LAG peers for this feature to work properly.

## RELATED DOCUMENTATION

| [Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up](#)

## forwarding-class (Gigabit Ethernet IQ Classifier)

### IN THIS SECTION

- Syntax | 885
- Hierarchy Level | 885
- Description | 885
- Options | 885
- Required Privilege Level | 886
- Release Information | 886

### Syntax

```
forwarding-class class-name {
 loss-priority (high | low);
}
```

### Hierarchy Level

```
[edit interfaces interface-name together-options ethernet-switch-profile
ethernet-policer-profile output-priority-map classifier premium]
```

### Description

For Gigabit Ethernet IQ interfaces only, define forwarding class name and option values.

### Options

*class-name*—Name of forwarding class.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[input-priority-map | 920](#)

[Junos OS Class of Service User Guide for Routing Devices](#)

## forwarding-mode (100-Gigabit Ethernet)

### IN THIS SECTION

- [Syntax | 886](#)
- [Hierarchy Level | 887](#)
- [Description | 887](#)
- [Required Privilege Level | 887](#)
- [Release Information | 887](#)

## Syntax

```
forwarding-mode {
 (sa-multicast | ...the following vlan-steering statement...);
 vlan-steering {
 vlan-rule (high-low | odd-even);
 }
}
```

```
}
}
```

## Hierarchy Level

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

## Description

Configure the interoperation mode for 100-Gigabit Ethernet PIC or the 100-Gigabit Ethernet MIC.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 10.4.

Statement introduced in Junos OS Release 12.1 for MX Series routers.

## RELATED DOCUMENTATION

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 231](#)

[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 | 238](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 239](#)

[sa-multicast \(100-Gigabit Ethernet\) | 1113](#)

[vlan-rule \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1225](#)

[vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1226](#)

## forwarding-mode (PTX Series Packet Transport Routers)

### IN THIS SECTION

- Syntax | 888
- Hierarchy Level | 888
- Description | 888
- Required Privilege Level | 888
- Release Information | 889

### Syntax

```
forwarding-mode {
 sa-multicast
}
```

### Hierarchy Level

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

### Description

Configure interoperability between 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP.

The remaining statement is explained separately. See [CLI Explorer](#).

### Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 12.1X48R4.

### RELATED DOCUMENTATION

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 242](#)

[Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | 238](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 239](#)

## frame-error

### IN THIS SECTION

- [Syntax | 889](#)
- [Hierarchy Level | 889](#)
- [Description | 890](#)
- [Options | 890](#)
- [Required Privilege Level | 890](#)
- [Release Information | 890](#)

### Syntax

```
frame-error count;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event link-event-rate],
```

```
[edit protocols oam link-fault-management interface interface-name event-
thresholds]
```

## Description

Threshold for sending frame error events or taking the action specified in the action profile.

A frame error is any frame error on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value within the window.

The window or period during which frame errors are counted is 5 seconds or multiples of it (with a maximum value of 1 minute). This window denotes the duration as intervals of 100 milliseconds, encoded as a 16-bit unsigned integer. This window is not configurable in Junos OS. According to the IEEE 802.3ah standard, the default value of the frame-errors window is 1 second. This window has a lower bound of 1 second and an upper bound of 1 minute.

## Options

*count*—Threshold count for frame error events.

- **Range:** 0 through 100

## 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 Release 8.4.

## RELATED DOCUMENTATION

*Configuring Threshold Values for Local Fault Events on an Interface*

*Configuring Threshold Values for Fault Events in an Action Profile*

## frame-period

### IN THIS SECTION

- [Syntax | 891](#)
- [Hierarchy Level | 891](#)
- [Description | 891](#)
- [Options | 891](#)
- [Required Privilege Level | 892](#)
- [Release Information | 892](#)

### Syntax

```
frame-period count;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event link-
event-rate],
[edit protocols oam link-fault-management interface interface-name event-
thresholds]
```

### Description

Threshold for sending frame period error events or taking the action specified in the action profile.

A frame error is any frame error on the underlying physical layer. The frame period threshold is reached when the number of frame errors reaches the configured value within the period window. The default period window is the number of minimum-size frames that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

### Options

*count*—Threshold count for frame period error events.

- **Range:** 0 through 100

## 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 Release 8.4.

## RELATED DOCUMENTATION

*Configuring Threshold Values for Local Fault Events on an Interface*

*Configuring Threshold Values for Fault Events in an Action Profile*

## frame-period-summary

### IN THIS SECTION

- [Syntax | 892](#)
- [Hierarchy Level | 893](#)
- [Description | 893](#)
- [Options | 893](#)
- [Required Privilege Level | 893](#)
- [Release Information | 893](#)

## Syntax

```
frame-period-summary count;
```

## Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event link-
event-rate],
[edit protocols oam link-fault-management interface interface-name event-
thresholds]
```

## Description

Threshold for sending frame period summary error events or taking the action specified in the action profile.

An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period window. The default window is 60 seconds. The window is not configurable.

## Options

*count*—Threshold count for frame period summary error events.

- **Range:** 0 through 100

## 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 Release 8.4.

## RELATED DOCUMENTATION

*Configuring Threshold Values for Local Fault Events on an Interface*

*Configuring Threshold Values for Fault Events in an Action Profile*

## framing (10-Gigabit Ethernet Interfaces)

### IN THIS SECTION

- [Syntax | 894](#)
- [Hierarchy Level | 894](#)
- [Description | 894](#)
- [Default | 895](#)
- [Options | 895](#)
- [Required Privilege Level | 895](#)
- [Release Information | 896](#)

### Syntax

```
framing (lan-phy | wan-phy);
precise-bandwidth;
```

### Hierarchy Level

```
[edit interfaces xe-fpc/pic/port]
```

```
[edit interfaces et-fpc/pic/port] (PTX Series Packet Transport Routers and MX
Series Routers)
```

### Description

For routers supporting the 10-Gigabit Ethernet interface, configure the framing format. WAN PHY mode is supported on MX240, MX480, MX960, T640, T1600, T4000, and PTX Series Packet Transport Routers 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).
- On PTX Series 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.
- In Junos OS Releases 17.4R2, 17.4R3, and later, on the following MPCs or routers, you cannot configure **wan-phy** mode at 10-Gbps, 40-Gbps, and 100-Gbps on a per-port basis:
  - MPC7E-10G, MPC7E-MRATE, MX2K-MPC8E, and MX2K-MPC9E
  - MPC10003
  - MX204 router
  - JNP10K-LC2101 MPC

**Default**

Operates in LAN PHY mode.

**Options**

**lan-phy**—10GBASE-R interface framing format that bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface.

**wan-phy**—10GBASE-W interface framing format that allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and SONET devices.

**precise-bandwidth**—Enables precise bandwidth for WAN-PHY interface framing format.

**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 Release 8.0.

The option **precise-bandwidth** introduced in Junos OS Release 19.3R1 for MX Series Routers.

## RELATED DOCUMENTATION

[Framing Overview | 202](#)

[Configuring SONET Options for 10-Gigabit Ethernet Interfaces](#)

## gigether-options

### IN THIS SECTION

- [Syntax | 896](#)
- [Hierarchy Level | 898](#)
- [Description | 898](#)
- [Required Privilege Level | 904](#)
- [Release Information | 904](#)

## Syntax

```
gigether-options {
 802.3ad {
 aex (primary | backup);
 lacp {
 port-priority;
 }
 }
 (asynchronous-notification | no-asynchronous-notification);
 (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-
online | local-interface-offline>;
 fec (gigether)
```

```

(flow-control | no-flow-control);
ignore-l3-incompletes;
(loopback | no-loopback);
loopback-remote
mpls {
 pop-all-labels {
 required-depth number;
 }
}
no-auto-mdix
source-address-filter {
 mac-address;
}
(source-filtering | no-source-filtering);
speed
ethernet-switch-profile {
 (mac-learn-enable | no-mac-learn-enable);
 tag-protocol-id [tpids];
 ethernet-policer-profile {
 input-priority-map {
 ieee802.1p premium [values];
 }
 output-priority-map {
 classifier {
 premium {
 forwarding-class class-name {
 loss-priority (high | low);
 }
 }
 }
 }
 policer cos-policer-name {
 aggregate {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 premium {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 }
 }
}

```

```
}
}
```

## Hierarchy Level

```
[edit interfaces interface-name]
```

## Description

Configure Gigabit Ethernet specific interface properties.

In Junos OS Evolved, when you configure **set interfaces *interface* gigether-options 802.3ad *ae name*** at the same time as you apply a second configuration to the same interface at the **[edit interfaces *interface*]** hierarchy, the second configuration will not take effect until the interface joins the aggregated Ethernet interface ***ae name***.

[Table 129 on page 899](#) shows the supported and unsupported platforms.

Table 129: Supported Platform Information

| Supported Platforms for gigheter-options                                                                                                                                                                                                                           | Supported Platforms for ether-options                                                                        | Notes                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| ACX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• ACX500</li> <li>• ACX710</li> <li>• ACX1000</li> <li>• ACX1100</li> <li>• ACX2100</li> <li>• ACX2200</li> <li>• ACX4000</li> <li>• ACX5448</li> <li>• ACX5048</li> <li>• ACX5096</li> </ul> | ACX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• ACX5048</li> <li>• ACX5096</li> </ul> | To configure storm control on the ACX Series routers, use <b>ether-options</b> . |

Table 129: Supported Platform Information (*Continued*)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                                                                                                        | Supported Platforms for <b>ether-options</b>                                                                                                                                                                                                                                           | Notes                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• EX2300</li> <li>• EX2300 Multigigabit</li> <li>• EX2300-C</li> <li>• EX3400</li> <li>• EX4300</li> <li>• EX4300 Multigigabit</li> <li>• EX4600</li> <li>• EX4650</li> <li>• EX9200</li> <li>• EX9250</li> </ul> | EX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• EX2300</li> <li>• EX2300 Multigigabit</li> <li>• EX2300-C</li> <li>• EX3400</li> <li>• EX4300</li> <li>• EX4300 Multigigabit</li> <li>• EX4600</li> <li>• EX4650</li> <li>• EX9200</li> <li>• EX9250</li> </ul> | EX Series Switches (Junos OS)<br><br>EX Series switches support both <b>ether-options</b> and <b>gigether-options</b> .<br><br>To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use <b>ether-options</b> .<br><br>To configure Forward Error Correction (FEC), use <b>gigether-options</b> . |

Table 129: Supported Platform Information (*Continued*)

| Supported Platforms for gigheter-options                                                                                                                                                                                                                                                                                                            | Supported Platforms for ether-options | Notes |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|-------|
| MX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• MX5</li> <li>• MX10</li> <li>• MX40</li> <li>• MX80</li> <li>• MX104</li> <li>• MX150</li> <li>• MX204</li> <li>• MX240</li> <li>• MX480</li> <li>• MX960</li> <li>• MX20008</li> <li>• MX2010</li> <li>• MX2020</li> <li>• MX10003</li> <li>• MX10008 and MX10016</li> </ul> | Not Supported                         | None  |

Table 129: Supported Platform Information (*Continued*)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                          | Supported Platforms for <b>ether-options</b>                                                                                                                       | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PTX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• PTX1000</li> <li>• PTX3000</li> <li>• PTX5000</li> <li>• PTX10001</li> <li>• PTX10002</li> <li>• PTX10008 and PTX10016</li> </ul> | PTX Series Routers (Junos OS) <ul style="list-style-type: none"> <li>• PTX1000</li> <li>• PTX10001</li> <li>• PTX10002</li> <li>• PTX10008 and PTX10016</li> </ul> | PTX Series Routers (Junos OS) <p>In Junos OS Release 17.3R3S7, PTX1000 Series routers support both <b>ether-options</b> and <b>gigether-options</b>.</p> <p>In Junos OS Releases 17.3R1, 17.4R1, and 17.4R2, PTX10000 Series routers support both <b>ether-options</b> and <b>gigether-options</b>.</p> <p>In Junos OS Release 18.1R1 and later, PTX Series routers support only <b>gigether-options</b> for your configuration and do not support <b>ether-options</b>. Before Junos OS Release 18.1R1, PTX Series support <b>ether-options</b>. See <a href="#">TSB17864</a> for additional information.</p> |
| PTX Series routers (Junos OS Evolved) <ul style="list-style-type: none"> <li>• PTX10003</li> <li>• PTX10008 and PTX10016</li> </ul>                                                                      | PTX Series routers (Junos OS Evolved) <ul style="list-style-type: none"> <li>• PTX10003</li> <li>• PTX10008 and PTX10016</li> </ul>                                | PTX Series routers (Junos OS Evolved) <p>Starting in Junos OS Evolved Release 20.1R1, PTX Series routers support <b>ether-options</b> only.</p> <p>Before Junos OS Evolved Release 20.1R1, PTX Series routers supported <b>gigether-options</b> only.</p>                                                                                                                                                                                                                                                                                                                                                      |

Table 129: Supported Platform Information (*Continued*)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Supported Platforms for <b>ether-options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Notes                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| QFX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• QFX5100 (48S)</li> <li>• QFX5100 (48T)</li> <li>• QFX5100 (24Q)</li> <li>• QFX5100 (96S)</li> <li>• QFX5110 (48S)</li> <li>• QFX5110 (32Q)</li> <li>• QFX5120 (48Y)</li> <li>• QFX5120 (48T)</li> <li>• QFX5120 (48ym)</li> <li>• QFX5120 (32C)</li> <li>• QFX5200 (48Y)</li> <li>• QFX5200 (32C)</li> <li>• QFX5210 (64C)</li> <li>• QFX5210 (64C-S)</li> <li>• QFX5220 (32CD)</li> <li>• QFX5220 (128C)</li> <li>• QFX10002</li> <li>• QFX10008 and QFX10016</li> </ul> | QFX Series Switches (Junos OS) <ul style="list-style-type: none"> <li>• QFX5100 (48S)</li> <li>• QFX5100 (48T)</li> <li>• QFX5100 (24Q)</li> <li>• QFX5100 (96S)</li> <li>• QFX5110 (48S)</li> <li>• QFX5110 (32Q)</li> <li>• QFX5120 (48Y)</li> <li>• QFX5120 (32C)</li> <li>• QFX5200 (48Y)</li> <li>• QFX5120 (48T)</li> <li>• QFX5120 (48ym)</li> <li>• QFX5200 (32C)</li> <li>• QFX5210 (64C)</li> <li>• QFX5210 (64C-S)</li> <li>• QFX5220 (32CD)</li> <li>• QFX5220 (128C)</li> <li>• QFX10002</li> <li>• QFX10008 and QFX10016</li> </ul> | QFX Series Switches (Junos OS) <p>QFX Series switches support both <b>ether-options</b> and <b>gigether-options</b>.</p> <p>To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use <b>ether-options</b>.</p> <p>To configure Forward Error Correction (FEC), use <b>gigether-options</b>.</p> |

Table 129: Supported Platform Information (*Continued*)

| Supported Platforms for <b>gigether-options</b>                                                                                                                                                               | Supported Platforms for <b>ether-options</b>                                                                                                                                                                  | Notes                                                                                                                                                                                                          |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Not Supported                                                                                                                                                                                                 | QFX Series Switches (Junos OS Evolved) <ul style="list-style-type: none"> <li>• QFX5200-32C-L</li> <li>• QFX5220-32CD</li> <li>• QFX5220-128C</li> </ul>                                                      | To configure auto-negotiation on management interfaces ( <b>re0:mgmt-0</b> ), use <b>ether-options</b> .                                                                                                       |
| SRX Series (Junos OS) <ul style="list-style-type: none"> <li>• SRX300</li> <li>• SRX550</li> <li>• SRX1500</li> <li>• SRX4100 and SRX4200</li> <li>• SRX4600</li> <li>• SRX5400</li> <li>• SRX5600</li> </ul> | SRX Series (Junos OS) <ul style="list-style-type: none"> <li>• SRX300</li> <li>• SRX550</li> <li>• SRX1500</li> <li>• SRX4100 and SRX4200</li> <li>• SRX4600</li> <li>• SRX5400</li> <li>• SRX5600</li> </ul> | SRX Series (Junos OS) <p>To configure gigabit-Ethernet interfaces (ge-), use <b>gigether-options</b>. To configure ethernet interfaces (et-) and fast ethernet interfaces (fe-), use <b>ether-options</b>.</p> |

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

### Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Ethernet Interfaces Overview | 2](#)

*ether-options*

## hash-key (Chassis LAG)

### IN THIS SECTION

- [Syntax | 905](#)
- [Hierarchy Level | 906](#)
- [Description | 906](#)
- [Options | 906](#)
- [Required Privilege Level | 906](#)
- [Release Information | 906](#)

### Syntax

```
hash-key {
 family {
 inet {
 layer-3;
 layer-4;
 symmetric-hash {
 complement;
 }
 }
 }
 multiservice {
 source-mac;
 destination-mac;
 payload {
 ip {
 layer-3 (source-ip-only | destination-ip-only);
 layer-4;
 }
 }
 }
}
```

```

 }
 }
}

```

## Hierarchy Level

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

## Description

(MX Series 5G Universal Routing Platforms only) Configure data used in a hash key for a PIC for symmetrical load balancing on an 802.3ad Link Aggregation Group.

## Options

**family**—Configure data used in a hash key for a protocol family. This statement has the following suboptions:

- **inet**—Configure data used in a hash key for the **inet** protocol family.
- **multiservice**—Configure data used in a hash key for the **multiservice** protocol family.

## 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 Release 9.6.

## RELATED DOCUMENTATION

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

## hold-time up

### IN THIS SECTION

- [Syntax | 907](#)
- [Hierarchy Level | 907](#)
- [Description | 907](#)
- [Options | 908](#)
- [Required Privilege Level | 908](#)
- [Release Information | 908](#)

### Syntax

```
hold-time up timer-value;
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options lacp],
```

### Description

Specifies the time period for which the Link Aggregation Control Protocol (LACP) maintains the state of a child (member) link as expired or default.

When a child link goes from the current state to the expired state, the LACP monitors the reception of protocol data units (PDUs) on the child link for the configured hold-up time interval and does not allow the child link to transition back to the current state. This configuration thus prevents excessive flapping of a child link on an aggregated Ethernet interface.

The configured hold-up timer value is applicable to all the child links within a link aggregated (LAG) interface. By default, this feature is disabled.

## Options

*timer-value*—Hold-up interval in seconds.

- **Range:** 1 to 6000 seconds

## 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 Release 14.2R3.

## RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#)

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

## ieee802.1p

### IN THIS SECTION

- [Syntax | 909](#)
- [Hierarchy Level | 909](#)
- [Description | 909](#)
- [Options | 909](#)
- [Required Privilege Level | 909](#)
- [Release Information | 909](#)

## Syntax

```
ieee802.1p premium [values];
```

## Hierarchy Level

```
[edit interfaces interface-name together-options ethernet-switch-profile
ethernet-policer-profile input-priority-map]
[edit interfaces interface-name ether-options ethernet-switch-profile ethernet-
policer-profile input-priority-map]
```

## Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, configure premium priority values for IEEE 802.1p input traffic.

## Options

*values*—Define IEEE 802.1p priority values to be treated as premium.

- **Range:** 0 through 7

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos Release 7.4.

## ignore-l3-incompletes

### IN THIS SECTION

- [Syntax | 910](#)
- [Hierarchy Level | 910](#)
- [Description | 910](#)
- [Required Privilege Level | 910](#)
- [Release Information | 910](#)

### Syntax

```
ignore-l3-incompletes;
```

### Hierarchy Level

```
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gether-options]
```

### Description

Ignore the counting of Layer 3 incomplete errors on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces.

### 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 Release 9.0.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

## RELATED DOCUMENTATION

| [Ignoring Layer 3 Incomplete Errors](#)

## inet (chassis)

### IN THIS SECTION

- [Syntax | 911](#)
- [Hierarchy Level | 911](#)
- [Description | 912](#)
- [Options | 912](#)
- [Required Privilege Level | 912](#)
- [Release Information | 912](#)

## Syntax

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

## Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key family]
```

## Description

(MX Series 5G Universal Routing Platforms only) Configure data used in a hash key for the **inet** protocol family when configuring PIC-level symmetrical load balancing on an 802.3ad Link Aggregation Group.

## Options

|                       |                                                                 |
|-----------------------|-----------------------------------------------------------------|
| <b>layer-3</b>        | Include Layer 3 IP data in the hash key.                        |
| <b>layer-4</b>        | Include Layer 4 IP data in the hash key.                        |
| <b>symmetric-hash</b> | Configure symmetric hash key with source and destination ports. |

## 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 Release 9.6.

## RELATED DOCUMENTATION

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

## ingress-policer-overhead

### IN THIS SECTION

- [Syntax | 913](#)
- [Hierarchy Level | 913](#)
- [Description | 913](#)
- [Options | 914](#)

- Required Privilege Level | 915
- Release Information | 915

## Syntax

```
ingress-policer-overhead bytes;
```

## Hierarchy Level

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

## Description

Add the configured number of bytes to the length of a packet entering the interface.

Configure a policer overhead to control the rate of traffic received on an interface. Use this feature to help prevent denial-of-service (DoS) attacks or to enforce traffic rates to conform to the service-level agreement (SLA). When you configure a policer overhead, the configured policer overhead value (bytes) is added to the length of the final Ethernet frame. This calculated length of frame is used to determine the policer or the rate-limiting action.

Traffic policing combines the configured policy bandwidth limits and the burst size to determine how to meter the incoming traffic. If you configure a policer overhead on an interface, Junos OS adds those bytes to the length of incoming Ethernet frames. This added overhead fills each frame closer to the burst size, allowing you to control the rate of traffic received on an interface.

You can configure the policer overhead to rate-limit queues and Layer 2 and Layer 3 policers, for standalone (SA) and high-availability (HA) deployments. The policer overhead and the shaping overhead can be configured simultaneously on an interface.

**NOTE:** vSRX supports policer overhead on Layer 3 policers only.

The policer overhead applies to all interfaces on the PIC. In the following example, Junos OS adds 10 bytes of overhead to all incoming Ethernet frames on ports ge-0/0/0 through ge-0/0/4.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 10
```

**NOTE:** vSRX only supports **fpc 0 pic 0**. When you commit the **ingress-policer-overhead** statement, the vSRX takes the PIC offline and then back online.

You need to craft the policer overhead size to match your network traffic. A value that is too low will have minimal impact on traffic bursts. A value that is too high will rate-limit too much of your incoming traffic.

In this example, the policer overhead of 255 bytes is configured for ge-0/0/0 through ge-0/0/4. The firewall policer is configured to discard traffic when the burst size is over 1500 bytes. This policer is applied to ge-0/0/0 and ge 0/0/1. Junos OS adds 255 bytes to every Ethernet frame that comes into the configured ports. If, during a burst of traffic, the combined length of incoming frames and the overhead bytes exceeds 1500 bytes, the policer starts to discard further incoming traffic.

```
set chassis fpc 0 pic 0 ingress-policer-overhead 255
set interfaces ge-0/0/0 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/0 unit 0 family inet address 10.9.1.2/24
set interfaces ge-0/0/1 unit 0 family inet policer input overhead_policer
set interfaces ge-0/0/1 unit 0 family inet address 10.9.2.2/24
set firewall policer overhead_policer if-exceeding bandwidth-limit 32k
set firewall policer overhead_policer if-exceeding burst-size-limit 1500
set firewall policer overhead_policer then discard
```

## Options

**bytes**—Number of bytes added to a frame entering an interface.

- **Range:** 0–255 bytes
- **Default:** 0

```
[edit chassis fpc 0 pic 0]
user@host# set ingress-policer-overhead 10;
```

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 11.1.

## RELATED DOCUMENTATION

*ingress-shaping-overhead*

*Policer Overhead to Account for Rate Shaping Overview*

*Example: Configuring Policer Overhead to Account for Rate Shaping*

*Configuring a Policer Overhead*

*CoS on Enhanced IQ2 PICs Overview*

## ingress-rate-limit

### IN THIS SECTION

- [Syntax | 915](#)
- [Hierarchy Level | 916](#)
- [Description | 916](#)
- [Options | 916](#)
- [Required Privilege Level | 916](#)
- [Release Information | 916](#)

## Syntax

```
ingress-rate-limit rate;
```

## Hierarchy Level

```
[edit interfaces interface-name fastether-options]
```

## Description

Perform port-based rate limiting on ingress traffic arriving on Fast Ethernet 8-port, 12-port, and 48-port PICs.

## Options

*rate*—Traffic rate, in megabits per second (Mbps).

- **Range:** 1 through 100 Mbps

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Configuring the Ingress Rate Limit](#) | 8

## inline

### IN THIS SECTION

- [Syntax](#) | 917
- [Hierarchy Level](#) | 917

- [Description | 917](#)
- [Default | 918](#)
- [Required Privilege Level | 918](#)
- [Release Information | 918](#)

## Syntax

```
inline;
```

## Hierarchy Level

```
[edit protocols lacp ppm]
```

## Description

(MX Series routers with MPC line cards only) To enable the inline Link Aggregation Control Protocol (LACP) PDU transmission processing.

This statement disables the default distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all Link Aggregation Control Protocol (LACP) PDU transmission processing inline. The inline option can be used in scenarios where the line card CPU is under heavy load and cannot schedule the PPM processing for LACP packets. PPM, by default, delegates the transmission of PDUs to the PPMAN process on the PFE/line card. But when the inline option is configured, it delegates the transmission of LACP PDUs even further away from the line card CPU and into the forwarding chipset.

For example, in a system with both MPCs and DPCs, upon configuration of **[protocols lacp ppm inline]**, the PDUs are sent inline on the MPCs and performed by periodic packet management (PPM) on DPCs.

**BEST PRACTICE:** We recommend to retain the default and disable distributed PPM or enable inline processing only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM or enable inline processing only if you have a compelling reason to disable it.

Refer [Disabling or Enabling Inline Periodic Packet Management for LACP Packets](#) for more details.

## Default

Distributed PPM processing is enabled for all packets that use PPM.

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 19.1R1.

## RELATED DOCUMENTATION

[centralized](#) | [821](#)

[Disabling or Enabling Inline Periodic Packet Management for LACP Packets](#)

*Configuring Link Aggregation*

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

## input-policer

### IN THIS SECTION

- [Syntax](#) | [919](#)
- [Hierarchy Level](#) | [919](#)
- [Description](#) | [919](#)
- [Options](#) | [919](#)
- [Usage Guidelines](#) | [919](#)
- [Required Privilege Level](#) | [919](#)
- [Release Information](#) | [919](#)

## Syntax

```
input-policer policer-name;
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer]
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number layer2-policer]
```

## Description

Apply a single-rate two-color policer to the Layer 2 input traffic at the logical interface. The **input-policer** and **input-three-color** statements are mutually exclusive.

## Options

***policer-name***—Name of the single-rate two-color policer that you define at the **[edit firewall]** hierarchy level.

## Usage Guidelines

See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

## 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 Release 8.2.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

---

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 271](#)

*input-three-color*

*layer2-policer*

*logical-interface-policer*

*output-policer*

*output-three-color*

## input-priority-map

### IN THIS SECTION

- [Syntax | 920](#)
- [Hierarchy Level | 920](#)
- [Description | 921](#)
- [Required Privilege Level | 921](#)
- [Release Information | 921](#)

### Syntax

```
input-priority-map {
 ieee802.1p premium [values];
}
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile]
[edit interfaces interface-name ether-options ethernet-switch-profile ethernet-
policer-profile]
```

## Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the input policer priority map to be applied to incoming frames on this interface.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[output-priority-map | 1055](#)

## input-three-color

### IN THIS SECTION

- [Syntax | 922](#)
- [Hierarchy Level | 922](#)
- [Description | 922](#)
- [Options | 922](#)
- [Usage Guidelines | 922](#)
- [Required Privilege Level | 922](#)
- [Release Information | 922](#)

## Syntax

```
input-three-color policer-name;
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer]
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number layer2-policer]
```

## Description

Apply a single-rate or two-rate three-color policer to the Layer 2 input traffic at the logical interface. The **input-three-color** and **input-policer** statements are mutually exclusive.

## Options

***policer-name***—Name of the single-rate or two-rate three-color policer.

## Usage Guidelines

See *Applying Layer 2 Policers to Gigabit Ethernet Interfaces*.

## 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 Release 8.2.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 271](#)

*input-policer*

*layer2-policer*

*logical-interface-policer*

*output-policer*

*output-three-color*

## input-vlan-map (Aggregated Ethernet)

### IN THIS SECTION

- [Syntax | 923](#)
- [Hierarchy Level | 923](#)
- [Description | 924](#)
- [Required Privilege Level | 924](#)
- [Release Information | 924](#)

### Syntax

```
input-vlan-map {
 (pop | push | swap);
 tag-protocol-id tpid;
 vlan-id number;
}
```

### Hierarchy Level

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

## Description

Define the rewrite profile to be applied to incoming frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet interfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.2.

Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

## RELATED DOCUMENTATION

[Stacking a VLAN Tag](#)

[output-vlan-map \(Aggregated Ethernet\) | 1058](#)

## interface (LLDP)

### IN THIS SECTION

- [Syntax | 925](#)
- [Hierarchy Level | 925](#)
- [Description | 925](#)
- [Options | 925](#)
- [Required Privilege Level | 926](#)
- [Release Information | 926](#)

## Syntax

```
interface (all | [interface-name-list]) {
 (disable | enable);
 power-negotiation <(disable | enable)>;
 (tlv-filter | tlv-select);
 trap-notification (disable | enable);
}
```

## Hierarchy Level

```
[edit protocols lldp],
[edit routing-instances routing-instance-name protocols lldp]
```

## Description

Configure Link Layer Discovery Protocol (LLDP) on all interfaces or on a particular interface.

**NOTE:** On MX Series and T Series routers, you run LLDP on a physical interface, such as **ge-1/0/0**, and not at the logical interface (unit) level.

Starting with Junos OS Release 14.2, on MX Series devices, you can also configure LLDP on management interfaces, such as **fxp** or **me**.

For information about interface names, see [Interface Naming Overview](#). For information about interface names for TX Matrix routers, see [TX Matrix Router Chassis and Interface Names](#). For information about FPC numbering on TX Matrix routers, see [Routing Matrix with a TX Matrix Router FPC Numbering](#).

For information about extended port names in the Junos Fusion technology, see *Understanding Junos Fusion Ports*.

## Options

- |                                             |                                                                          |
|---------------------------------------------|--------------------------------------------------------------------------|
| <b>(all   [<i>interface-name-list</i>])</b> | Configure LLDP on all interfaces or on one or more interfaces.           |
| <b>(disable   enable)</b>                   | Disable or enable LLDP on all interfaces or on the specified interfaces. |

- **Default:** Disable

**power-negotiation**  
<(disable |  
enable)>

(EX, QFX Series only) Configure LLDP power negotiation, which negotiates with Power over Ethernet (PoE) powered devices to allocate power.

You must also configure the **management class** statement at the [edit poe] hierarchy level to activate LLDP power negotiation.

- **Values:** Configure one of the following:
  - disable—Disable LLDP power negotiation.
  - enable—Enable LLDP power negotiation.

**trap-notification**  
(disable | enable)

Disables or enables the LLDP and physical topology SNMP traps for the specific interface or all the interfaces.

- **Values:** Configure one of the following:
  - disable—Disable the LLDP and physical topology SNMP trap notifications.
  - enable—Enable the LLDP and physical topology SNMP trap notifications.
- **Default:** disable

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

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.0.

**power-negotiation** introduced in Junos OS Release 12.2 for EX and QFX Series switches.

**trap-notification** introduced in Junos OS Release 15.1R7-S3 for EX3300, EX4200, EX4500, EX4550, EX6200, EX8200 switches.

## RELATED DOCUMENTATION

[Configuring LLDP \(CLI Procedure\)](#)

[Configuring LLDP](#)

*Configuring PoE Interfaces on EX Series Switches*

## interface (OAM Link-Fault Management)

### IN THIS SECTION

- [Syntax | 927](#)
- [Hierarchy Level | 928](#)
- [Description | 928](#)
- [Options | 928](#)
- [Required Privilege Level | 928](#)
- [Release Information | 928](#)

### Syntax

```
interface interface-name {
 apply-action-profile profile-name;
 link-discovery (active | passive);
 pdu-interval interval;
 pdu-threshold threshold-value;
 remote-loopback;
 event-thresholds {
 frame-error count;
 frame-period count;
 frame-period-summary count;
 symbol-period count;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
}
```

```
}
}
```

## Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

## Description

For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.3ah Operation, Administration, and Management (OAM) support.

## Options

**interface** *interface-name*—Interface to be enabled for IEEE 802.3ah link fault management OAM support.

The remaining statements are described separately.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.2.

## RELATED DOCUMENTATION

| [Enabling IEEE 802.3ah OAM Support](#)

## interface-group

### IN THIS SECTION

- [Syntax | 929](#)
- [Hierarchy Level | 929](#)
- [Description | 929](#)
- [Options | 929](#)
- [Required Privilege Level | 930](#)
- [Release Information | 930](#)

### Syntax

```
interface-group {
 interface-device-name
 unit-list
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association ma-name mep mep-id remote-mep mep-id]
```

### Description

Mark the interface group down for the action profile configured with the action **interface-group-down**. Provides information for the interface-group on which the configured action will be taken when the configured event occur for a specific remote MEP ID.

### Options

***interface-device name*** Name of the interface device. Only Ethernet devices are allowed. The device interface name includes **ge, ae, xe** and **et** .

**unit-list** One or more logical interface unit numbers.

- **Range:** A string in the range <0-16385> or <0-16385>-<0-16385>. For example, **unit-list**[12 23-33 44]

## 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 Release 18.1R1.

## RELATED DOCUMENTATION

*Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces*

[interface-group-down](#) | 930

## interface-group-down

### IN THIS SECTION

- [Syntax](#) | 930
- [Hierarchy Level](#) | 931
- [Description](#) | 931
- [Required Privilege Level](#) | 931
- [Release Information](#) | 931

## Syntax

```
interface-group-down
```

## Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action]
```

## Description

Mark the interface group down.

## 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 Release 18.1R1.

## RELATED DOCUMENTATION

*Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces*

[interface-group](#) | 929

## interface-none

### IN THIS SECTION

- [Syntax](#) | 932
- [Hierarchy Level](#) | 932
- [Description](#) | 932
- [Required Privilege Level](#) | 932

## Syntax

```
interface-none;
```

## Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name east-interface]
```

```
[edit protocols protection-group ethernet-ring ring-name west-interface]
```

## Description

Designates port as not used for Ethernet ring protection.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Ring Protection Switching Overview](#)

[Ethernet Ring Protection Using Ring Instances for Load Balancing](#)

[Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)

[Configuring Ethernet Ring Protection Switching on Switches \(CLI Procedure\)](#)

## isolated-vlan (MX Series)

### IN THIS SECTION

● [Syntax | 933](#)

- [Hierarchy Level | 933](#)
- [Description | 933](#)
- [Options | 934](#)
- [Required Privilege Level | 934](#)
- [Release Information | 934](#)

## Syntax

```
isolated-vlan vlan-id;
```

## Hierarchy Level

```
[edit bridge-domains bridge-domain-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains bridge-domain-name bridge-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name bridge-domains bridge-domain-name ,
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name],
```

## Description

Configure the specified isolated VLAN to be a secondary VLAN of the specified primary VLAN. An isolated VLAN receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.

**NOTE:** When you specify this configuration statement, the VLAN ID of a logical interface that you associate with a bridge domain that matches with the VLAN ID that you specify using the `isolated-vlan` state is treated as an isolated port.

## Options

*vlan-id* Individual VLAN IDs separated by a space.

## Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 14.2.

## iwf-params (tdm-options)

### IN THIS SECTION

- [Syntax | 934](#)
- [Hierarchy Level | 935](#)
- [Description | 935](#)
- [Default | 935](#)
- [Options | 935](#)
- [Required Privilege Level | 935](#)
- [Release Information | 935](#)

## Syntax

```
iwf-params {
 decap-ecid value;
 encap-ecid value;
}
```

## Hierarchy Level

```
[edit interfaces interface-name tdm-options]
```

## Description

Configure TDM interworking functionality (IWF) parameters. You can configure these parameters to support multiple streams.

## Default

If you do not specify an emulation circuit id, the default emulation circuit id for decapsulation and encapsulation is 0. You cannot configure multiple streams if the emulation circuit id is 0.

## Options

**decap-ecid *value*** Emulation circuit id for de-encapsulation. Possible values: 0 through 1048575.

**encap-ecid *value*** Emulation circuit id for encapsulation. Possible values: 0 through 1048575.

## 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 Release 19.4.

## RELATED DOCUMENTATION

[show interfaces smart-sfp-statistics | 1719](#)

[show interfaces smart-sfp-defects | 1705](#)

## lacp (802.3ad)

### IN THIS SECTION

- [Syntax | 936](#)
- [Hierarchy Level | 936](#)
- [Description | 936](#)
- [Options | 936](#)
- [Required Privilege Level | 937](#)
- [Release Information | 937](#)

### Syntax

```
lacp {
 port-priority port-priority;
}
```

### Hierarchy Level

```
[edit interfaces interface-name fastether-options 802.3ad],
[edit interfaces interface-name ggether-options 802.3ad]
```

### Description

Configure the Link Aggregation Control Protocol (LACP) port priority for Ethernet interfaces.

### Options

*port-priority*—Priority for being elected as the active port to collect and distribute traffic. A smaller value indicates a higher priority for selection.

- **Range:** 0 through 65,535
- **Default:** 127

## 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 Release 9.3.

## RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#)

[port-priority](#) | [1078](#)

## lACP (Aggregated Ethernet)

### IN THIS SECTION

- [Syntax \(NFX Series\)](#) | [937](#)
- [Syntax \(EX Series\)](#) | [938](#)
- [Hierarchy Level \(EX Series\)](#) | [938](#)
- [Hierarchy Level \(NFX Series\)](#) | [938](#)
- [Description](#) | [939](#)
- [Default](#) | [939](#)
- [Options](#) | [939](#)
- [Required Privilege Level](#) | [940](#)
- [Release Information](#) | [940](#)

## Syntax (NFX Series)

```
lACP (active | passive) {
 admin-key key;
```

```

fast-failover;
link-protection {
 disable;
 (revertive | non-revertive);
}
periodic interval
system-ID mac-address;
system-priority priority;
force-up;
}

```

## Syntax (EX Series)

```

lacp {
 (active | passive);
 admin-key key;
 accept-data;
 fast-failover;

 link-protection {
 disable;
 (revertive | non-revertive);
 }
 periodic interval;
 system-id mac-address;
 system-priority priority;
}

```

## Hierarchy Level (EX Series)

```

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

```

## Hierarchy Level (NFX Series)

```

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

```

## Description

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

For EX Series, when you configure the **accept-data** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

**NOTE:** When you configure the **accept-data** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

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

## Default

If you do not specify LACP as either **active** or **passive**, LACP remains passive.

## Options

**active**                      Initiate transmission of LACP packets.

**admin-key  
number**                      Specify an administrative key for the router or switch.

**NOTE:** You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

- fast-failover** Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.
- passive** Respond to LACP packets.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.0.

## RELATED DOCUMENTATION

*Configuring Link Aggregation*

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

*Understanding Aggregated Ethernet Interfaces and LACP for Switches*

[Configuring LACP for Aggregated Ethernet Interfaces](#)

## lACP

### IN THIS SECTION

- [Syntax | 941](#)
- [Hierarchy Level | 941](#)
- [Description | 941](#)
- [Options | 941](#)
- [Required Privilege Level | 941](#)
- [Release Information | 941](#)

## Syntax

```
lACP {
 link-protection {
 non-revertive;
 }
 system-priority priority;
}
```

## Hierarchy Level

```
[edit chassis aggregated-devices ethernet]
```

## Description

For aggregated Ethernet interfaces only, configure Link Aggregation Control Protocol (LACP) parameters at the global level for use by LACP at the interface level.

## Options

The statements are described separately.

## 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 Release 9.3.

## RELATED DOCUMENTATION

| [Configuring Junos OS for Supporting Aggregated Devices](#)

## lACP (Protocols)

### IN THIS SECTION

- [Syntax | 942](#)
- [Hierarchy Level | 942](#)
- [Description | 942](#)
- [Default | 943](#)
- [Options | 943](#)
- [Required Privilege Level | 943](#)
- [Release Information | 943](#)

### Syntax

```
lACP {
 traceoptions {
 file <filename> <files number> <size size> <world-readable | no-world-
readable>;
 flag flag;
 no-remote-trace;
 }
 fast-hello-issu;
 ppm (Ethernet Switching) centralized;
}
```

### Hierarchy Level

```
[edit protocols]
```

### Description

On MX and T Series routers, you can specify periodic packet management (PPM) as centralized. By default, the PPM is distributed.

MX Series routers support Link Aggregation Control Protocol (LACP) with fast hellos during unified ISSU. This support is disabled by default. You must enable the **fast-hello-issu** option on the main router and on the peer routers before starting unified ISSU. Note that the peer router must also be an MX Series router for this functionality to work.

## Default

Distributed PPM processing is enabled for all packets that use PPM.

## Options

**ppm**—Set PPM to centralized.

**fast-hello-issu**—Enable LACP with fast hellos during unified ISSU.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.3.

The **ppm centralized** option introduced in Junos OS Release 9.4.

The **fast-hello-issu** option introduced in Junos OS Release 14.1.

## RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#)

[Unified ISSU System Requirements](#)

# lane

## IN THIS SECTION

- [Syntax | 944](#)
- [Hierarchy Level | 944](#)
- [Description | 945](#)
- [Options | 945](#)
- [Required Privilege Level | 947](#)
- [Release Information | 947](#)

## Syntax

```
lane [lane-number] {
 rx-cdr (off | on);
 rx-output (off | on);
 rx-output-amplitude (1200mV | 400mV | 600mV | 800mV);
 rx-output-emphasis rx-output-e;
 rx-rate-select rx-rate-select;
 rx-squelch (off | on);
 tx-cdr (off| no);
 tx-input-equalization-adaptive (off | on);
 tx-input-equalization-manual tx-input-manual;
 tx-laser (off | on);
 tx-rate-select tx-rate-select;
 tx-squelch (off | on);
}
```

## Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

## Description

Configure the optical module parameters of QSFP (Quad Small Form-Factor Pluggable) modules for a specific media lane number. The following QSFP modules are supported:

- QSFP+ : 4 host lanes, supports 10g per lane, total bandwidth of 40g
- QSFP28 : 4 host lanes, supports 25g per lane, total bandwidth of 100g
- QSFP56 : 4 host lanes, supports 50g per lane, total bandwidth of 200g
- QSFP28-DD : 8 host lanes, supports 25g per lane, total bandwidth of 200g
- QSFP56-DD : 8 host lanes, supports 50g per lane, total bandwidth of 400g

The following specifications standardize how the host device interacts with the QSFP optical modules:

- The 4-lane QSFP modules adhere to SFF-8636 management interface specification.
- The 8-lane QSFP-DD modules adhere to CMIS (Common Management Interface Specification).

You can configure optical module parameters for QSFP28 and QSFP28-DD optical modules for Juniper qualified optics as well as third-party optics from Junos OS. You cannot configure optical module parameters for QSFP copper DAC (direct attach cables) or QSFP load modules.

## Options

|                            |                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>lane-number</b>         | Specifies the media lane number. <ul style="list-style-type: none"> <li>• Range: 0 through 7</li> </ul>                                                                                                                                    |
| <b>rx-cdr</b>              | Controls receive clock and data recovery block of the optics. <ul style="list-style-type: none"> <li>• <b>Off</b>—Receive clock and data recovery is disabled.</li> <li>• <b>On</b>—Receive clock and data recovery is enabled.</li> </ul> |
| <b>rx-output</b>           | Controls receive output from the optics towards the host device. <ul style="list-style-type: none"> <li>• <b>Off</b>—Receive output for optics is disabled.</li> <li>• <b>On</b>—Receive output for optics is enabled.</li> </ul>          |
| <b>rx-output-amplitude</b> | Controls the amplitude of receive output for which equalization is to be applied. <ul style="list-style-type: none"> <li>• <b>1200mV</b>—No equalization for 600mV - 1200mV amplitude.</li> </ul>                                          |

- **400mV**—No equalization for 100mV - 400mV amplitude.
  - **600mV**—No equalization for 300mV - 600mV amplitude.
  - **800mV**—No equalization for 400mV - 800mV amplitude.
- rx-output-emphasis**      Equalization value (in dB) for receive output signal of the optics.
- Range: 0 through 10
- rx-rate-select**      Controls baud rate for receive output signal.
- Range: 0 through 3
- rx-squelch**      Squelch the receive output signal of optics.
- **Off**—Receive output signal for the optics is disabled.
  - **On**—Receive output signal for the optics is enabled.
- tx-cdr**      Controls transmit clock and data recovery block of the optics.
- **Off**—Transmit clock and data recovery is disabled.
  - **On**—Transmit clock and data recovery is enabled.
- tx-input-equalization-adaptive**      Controls adaptive equalization for optics transmit input.
- **Off**—Adaptive equalization is disabled.
  - **On**—Adaptive equalization is enabled.
- tx-input-equalization-manual**      Equalization value (in dB) for optics transmit input.
- Range: 0 through 15
- tx-laser**      Controls transmit laser block of the optics.
- **Off**—Transmit laser is disabled.
  - **On**—Transmit laser is enabled.
- tx-rate-select**      Controls baud rate for transmit input signal.
- Range: 0 through 3
- tx-squelch**      Squelch the transmit input signal of the optics.
- **Off**—Transmit input signal for the optics is disabled.

- **On**—Transmit input signal for the optics is enabled.

## 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 21.1R1.

## RELATED DOCUMENTATION

[optics-options](#) | [1048](#)

[lane-all](#) | [947](#)

## lane-all

### IN THIS SECTION

- [Syntax](#) | [947](#)
- [Hierarchy Level](#) | [948](#)
- [Description](#) | [948](#)
- [Options](#) | [949](#)
- [Required Privilege Level](#) | [950](#)
- [Release Information](#) | [950](#)

## Syntax

```
lane-all {
 rx-cdr (off | on);
 rx-output (off | on);
}
```

```

rx-output-amplitude (1200mV | 400mV | 600mV | 800mV);
rx-output-emphasis rx-output-e;
rx-rate-select rx-rate-select;
rx-squelch (off | on);
tx-cdr (off | no);
tx-input-equalization-adaptive (off | on);
tx-input-equalization-manual tx-input-manual;
tx-laser (off | on);
tx-rate-select tx-rate-select;
tx-squelch (off | on);
}

```

## Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

## Description

Configure the optical module parameters of QSFP (Quad Small Form-Factor Pluggable) modules for all the media lanes. The following QSFP modules are supported:

- QSFP+ : 4 host lanes, supports 10g per lane, total bandwidth of 40g
- QSFP28 : 4 host lanes, supports 25g per lane, total bandwidth of 100g
- QSFP56 : 4 host lanes, supports 50g per lane, total bandwidth of 200g
- QSFP28-DD : 8 host lanes, supports 25g per lane, total bandwidth of 200g
- QSFP56-DD : 8 host lanes, supports 50g per lane, total bandwidth of 400g

The following specifications standardize how the host device interacts with the QSFP optical modules:

- The 4-lane QSFP modules adhere to SFF-8636 management interface specification.
- The 8-lane QSFP-DD modules adhere to CMIS (Common Management Interface Specification).

You can configure optical module parameters for QSFP28 and QSFP28-DD optical modules for Juniper qualified optics as well as third-party optics from Junos OS. You cannot configure optical module parameters for QSFP copper DAC (direct attach cables) or QSFP load modules.

## Options

- rx-cdr** Controls receive clock and data recovery block of the optics.
- **Off**—Receive clock and data recovery is disabled.
  - **On**—Receive clock and data recovery is enabled.
- rx-output** Controls receive output from the optics towards the host device.
- **Off**—Output for optics is disabled.
  - **On**—Output for optics is enabled.
- rx-output-amplitude** Controls the amplitude of receive output for which equalization is to be applied.
- **1200mV**—No equalization for 600mV - 1200mV amplitude.
  - **400mV**—No equalization for 100mV - 400mV amplitude.
  - **600mV**—No equalization for 300mV - 600mV amplitude.
  - **800mV**—No equalization for 400mV - 800mV amplitude.
- rx-output-emphasis** Equalization value (in dB) for receive output signal of the optics.
- Range: 0 through 15
- rx-rate-select** Controls baud rate for receive output signal.
- Range: 0 through 3
- rx-squelch** Squelches the receive output signal of optics.
- **Off**—Receive output signal for the optics is disabled.
  - **On**—Receive output signal for the optics is enabled.
- tx-cdr** Control transmit clock and data recovery block of the optics.
- **Off**—Transmit clock and data recovery is disabled.
  - **On**—Transmit clock and data recovery is enabled.
- tx-input-equalization-adaptive** Controls adaptive equalization for optics transmit input.
- **Off**—Transmit adaptive equalization is disabled.
  - **On**—Transmit adaptive equalization is enabled.

- tx-input-equalization-manual** Specifies the manual transmit equalization status of the optics.
- Range: 0 through 15
- tx-laser** Specifies the transmit laser status of the optics.
- **Off**—Transmit laser is disabled.
  - **On**—Transmit laser is enabled.
- tx-rate-select** Specifies if the optics can transmit rate selection.
- Range: 0 through 3
- tx-squelch** Squelches the transmit input signal of the optics.
- **Off**—Transmit input signal for the optics is disabled.
  - **On**—Transmit input signal for the optics is enabled.

### 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 21.1R1.

### RELATED DOCUMENTATION

[optics-options | 1048](#)

[lane | 944](#)

## layer2-policer

### IN THIS SECTION

[Syntax | 951](#)

- [Hierarchy Level | 951](#)
- [Description | 951](#)
- [Options | 952](#)
- [Required Privilege Level | 952](#)
- [Release Information | 952](#)

## Syntax

```
layer2-policer {
 input-policer policer-name;
 input-three-color policer-name;
 output-policer policer-name;
 output-three-color policer-name;
}
```

## Hierarchy Level

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

## Description

For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series, MX Series, and T Series routers, and for aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on EX Series switches, apply Layer 2 logical interface policers. The following policers are supported:

- Two-color
- Single-rate tricolor marking (srTCM)
- Two-rate tricolor marking (trTCM)

Two-color and tricolor policers are configured at the **[edit firewall]** hierarchy level.

## Options

**input-policer** *policer-name*—Two-color input policer to associate with the interface. This statement is mutually exclusive with the **input-three-color** statement.

**input-three-color** *policer-name*—Tricolor input policer to associate with the interface. This statement is mutually exclusive with the **input-policer** statement.

**output-policer** *policer-name*—Two-color output policer to associate with the interface. This statement is mutually exclusive with the **output-three-color** statement.

**output-three-color** *policer-name*—Tricolor output policer to associate with the interface. This statement is mutually exclusive with the **output-policer** statement.

## 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 Release 8.2.

## RELATED DOCUMENTATION

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 279](#)

## link-adjacency-loss

### IN THIS SECTION

- [Syntax | 953](#)
- [Hierarchy Level | 953](#)
- [Description | 953](#)
- [Required Privilege Level | 953](#)

## Syntax

```
link-adjacency-loss;
```

## Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event]
```

## Description

Loss of adjacency with IEEE 802.3ah link-fault management peer event. When included, the loss-of-adjacency event triggers the action specified under the **action** statement.

## 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 Release 8.5.

## RELATED DOCUMENTATION

| *Monitoring the Loss of Link Adjacency*

## link-discovery

### IN THIS SECTION

- [Syntax | 954](#)
- [Hierarchy Level | 954](#)
- [Description | 954](#)
- [Options | 954](#)
- [Required Privilege Level | 955](#)
- [Release Information | 955](#)

### Syntax

```
link-discovery (active | passive);
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

### Description

For Ethernet interfaces on EX Series switches, and M320, M120, MX Series, and T Series routers, specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

### Options

**(active | passive)**—Passive or active mode. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. Once the discovery process is initiated, both sides participate in discovery.

## 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 Release 8.2.

## RELATED DOCUMENTATION

| *Configuring Link Discovery*

## link-degrade-monitor

### IN THIS SECTION

- [Syntax | 955](#)
- [Hierarchy Level | 956](#)
- [Description | 956](#)
- [Options | 957](#)
- [Required Privilege Level | 957](#)
- [Release Information | 957](#)

## Syntax

```
link-degrade-monitor {
 (link-degrade-monitor-enable | no-link-degrade-monitor-enable);
 actions media-based;
 recovery {
 (auto | manual);
 timer timer;
 }
}
```

```

thresholds {
 clear clear-value;
 interval interval-value;
 set set-value;
 warning-clear warning-clear-value;
 warning-set warning-set-value;
}
}

```

## Hierarchy Level

```
[edit interfaces interfaces-name]
```

## Description

Configure link degrade monitoring on an interface and when a link degrade event is detected, specify the corrective action to be triggered. When configured, the feature monitors the quality of physical links on Ethernet interfaces (10-Gigabit, 25-Gigabit, 40-Gigabit, 50-Gigabit, 100-Gigabit, 200-Gigabit and 400-Gigabit) and triggers the user-configured action when the link's bit error rate (BER) value breaches the configured threshold. This feature can detect a BER value as low as  $10^{-15}$  to  $10^{-1}$ .

Starting in Junos OS Evolved Release 20.2R1, you can monitor physical link degradation based on Bit Error Rate (BER) on 10-Gigabit, 25-Gigabit, 40-Gigabit, 50-Gigabit, 100-Gigabit, 200-Gigabit and 400-Gigabit Ethernet interfaces on PTX1001-36MR. The guidelines to configure **link-degrade-monitor** on Gigabit Ethernet interfaces are:

- Whenever there is a local degrade event, in any of the interfaces, a **LOCAL-DEGRADE** flag is set and a **REMOTE-DEGRADE** flag is advertised on the peer interface of the connected router. When you enable the **link-degrade** option, Junos OS will maintain statistics such as, number of times the link degrade event has occurred and if the link is degraded, then seconds elapsed since the link degrade event. You can view these details by using the ["show interfaces " on page 1463](#) command. This command also shows the **Link Degrade Status** (if the link degrade **Defect**) is active or cleared. See ["show interfaces " on page 1463](#) for output information.
- On 400-Gigabit Ethernet interfaces:
  - The permissible range of BER is between 1E-5 to 1E-10. If you set any threshold value beyond the permissible range the value will be auto adjusted between 1E-5 to 1E-10. If you configure the set threshold value higher than 1E-5 (i.e between 1E-1 to 1E-4), the value will be adjusted to 1E-5. If you configure the clear threshold value lesser than 1E-10 (i.e between 1E-11 and 1E-15), the value will be adjusted to 1E-10.

- The option **threshold intervals** and **recover timer** are ignored. The error rates for 400-Gigabit Ethernet are always calculated per second. As soon as the BER per second becomes lower than clear threshold error rate, the link will be considered as recovered from degradation.
- The bit errors detected before applying FEC (Pre-FEC) correction is considered for determining the degrade state. The **set** and **clear** thresholds apply for pre FEC bit errors. For information FEC support see, [Port Speed on PTX10001-36MR Router Overview](#).

## Options

|                                       |                                                                                                                                                                                                                                            |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>link-degrade-monitor-enable</b>    | Enable link degrade monitoring on an interface.                                                                                                                                                                                            |
| <b>no-link-degrade-monitor-enable</b> | Disable link degrade monitoring on an interface.                                                                                                                                                                                           |
| <b>media-based</b>                    | Action to be taken when a link degrade event is detected. A media-based action brings down the physical link at both local and remote ends of the interface, and stops BER monitoring at the local end until an autorecovery is triggered. |

The remaining statements are described separately.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 15.1.

The options **link-degrade-monitor-enable** and **no-link-degrade-monitor-enable** are introduced in Junos OS Evolved Release 20.1R1.

## RELATED DOCUMENTATION

[Link Degrade Monitoring Overview | 600](#)

[show interfaces | 1463](#)

[thresholds | 1179](#)

[recovery | 1101](#)

[request interface link-degrade-recover | 1243](#)

*Common Output Fields Description*

## link-down

### IN THIS SECTION

- [Syntax | 958](#)
- [Hierarchy Level | 958](#)
- [Description | 958](#)
- [Required Privilege Level | 958](#)
- [Release Information | 958](#)

### Syntax

```
link-down;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

### Description

Mark the interface down for transit traffic.

### 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 Release 8.5.

## RELATED DOCUMENTATION

| *Specifying the Actions to Be Taken for Link-Fault Management Events*

## link-event-rate

### IN THIS SECTION

- [Syntax | 959](#)
- [Hierarchy Level | 959](#)
- [Description | 959](#)
- [Required Privilege Level | 960](#)
- [Release Information | 960](#)

### Syntax

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

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event]
```

### Description

Configure the number of link-fault management events per second.

## 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 Release 8.5.

## RELATED DOCUMENTATION

| *Configuring Threshold Values for Fault Events in an Action Profile*

## link-fault-management

### IN THIS SECTION

- [Syntax | 960](#)
- [Hierarchy Level | 961](#)
- [Description | 961](#)
- [Required Privilege Level | 962](#)
- [Release Information | 962](#)

## Syntax

```
link-fault-management {
 action-profile profile-name {
 action {
 link-down;
 send-critical-event;
 syslog;
 }
 event {
```

```

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

```

## Hierarchy Level

```
[edit protocols oam ethernet]
```

## Description

For Ethernet interfaces on M320, M120, MX Series, and T Series routers and EX Series switches, specify fault signaling and detection for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.2.

## RELATED DOCUMENTATION

| [Enabling IEEE 802.3ah OAM Support](#)

## link-mode

### IN THIS SECTION

- [Syntax | 962](#)
- [Hierarchy Level | 963](#)
- [Description | 963](#)
- [Options | 963](#)
- [Required Privilege Level | 964](#)
- [Release Information | 964](#)

## Syntax

```
link-mode mode (automatic | full-duplex | half-duplex);
```

## Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name ether-options],
[edit interfaces ge-pim/0/0 switch-options switch-port port-number]
```

## Description

Set the device's link connection characteristic.

## Options

*mode*—Link characteristics:

- **automatic**—Link mode is negotiated. This is the default for EX Series switches.
- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.
- **Default:** Fast Ethernet interfaces can operate in either full-duplex or half-duplex mode. The router's or switch's management Ethernet interface, **fxp0** or **em0**, and the built-in Fast Ethernet interfaces on the FIC (M7i router) autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.

**NOTE:** On EX Series switches, if **no-auto-negotiation** is specified in `[edit interfaces interface-name ether-options]`, you can select only **full-duplex** or **half-duplex**. If **auto-negotiation** is specified, you can select any mode.

**NOTE:**

- Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.
- Starting with Junos OS release 16.1R7 and later, the **link-mode** configuration is not supported on 10-Gigabit Ethernet Interfaces.

- Starting in Junos OS release 18.4R1, half-duplex mode is supported on SRX340 and SRX345 devices.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

## RELATED DOCUMENTATION

[Configuring the Link Characteristics on Ethernet Interfaces](#)

[Understanding Management Ethernet Interfaces](#)

[Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

[Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support](#)

## link-protection

### IN THIS SECTION

- [Syntax | 965](#)
- [Hierarchy Level | 965](#)
- [Description | 965](#)
- [Options | 965](#)
- [Required Privilege Level | 965](#)
- [Release Information | 966](#)

## Syntax

```
link-protection {
 disable;
 (revertive |non-revertive);
}
```

## Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
[edit interfaces aex aggregated-ether-options lacp]
```

## Description

On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the `[edit interfaces ge-fpc/pic/port gigerther-options 802.3ad aex]` hierarchy level or the `[edit interfaces fe-fpc/pic/port fastether-options 802.3ad aex]` hierarchy level.

On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.

For Junos OS link protection, specify **link-protection** at the following hierarchy levels:

- `[edit interfaces ge-fpc/pic/port ether-options 802.3ad aex]`
- `[edit interfaces xe-fpc/pic/port ether-options 802.3ad aex]` hierarchy level or at the `[edit interfaces xe-fpc/pic/port ether-options 802.3ad aex]` hierarchy level.

To disable link protection, use the **delete interface ae aggregate-ether-options link-protection** statement at the `[edit interfaces aex aggregated-ether-options]` hierarchy level or the `[edit interfaces aex aggregated-ether-options lacp]` hierarchy level.

## Options

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.3.

Support for **disable**, **revertive**, and **non-revertive** statements added in Junos OS Release 9.3.

## RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Link Protection | 106](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## link-protection (non-LACP)

### IN THIS SECTION

- [Syntax | 966](#)
- [Hierarchy Level | 967](#)
- [Description | 967](#)
- [Options | 967](#)
- [Required Privilege Level | 967](#)
- [Release Information | 967](#)

## Syntax

```
link-protection {
 link-protection-revertive;
}
```

## Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

## Description

User can specify the [link-protection-revertive](#) statement in the link protection configuration at the aggregated Ethernet interface level to set revertive mode. In revertive mode, adding a higher-priority link to the aggregated Ethernet bundle results in recalculation of the priorities and traffic will switch from the currently active link to the newly added, higher-priority link. Recalculation of priorities is performed only while link event such as addition\deletion and UP/Down operation on link, that is, configuration of this option will not result in any recalculation immediately until next link-event occurs.

In addition to enabling static link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the `[edit interfaces ge-fpc/pic/port gigheter-options 802.3ad aex]` hierarchy level or the `[edit interfaces fe-fpc/pic/port fastether-options 802.3ad aex]` hierarchy level.

For static link protection, specify **link-protection** at the following hierarchy levels:

- `[edit interfaces ge-fpc/pic/port ether-options 802.3ad aex]`
- `[edit interfaces xe-fpc/pic/port ether-options 802.3ad aex]` hierarchy level or at the `[edit interfaces xe-fpc/pic/port ether-options 802.3ad aex]` hierarchy level.

To disable static link protection, use the **delete interface ae aggregate-ether-options link-protection** statement at the `[edit interfaces aex aggregated-ether-options]` hierarchy level.

## Options

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 17.3R1.

## RELATED DOCUMENTATION

| [Configuring Aggregated Ethernet Link Protection](#) | 106

## link-protection (Protocols LACP)

### IN THIS SECTION

- [Syntax](#) | 968
- [Hierarchy Level](#) | 968
- [Description](#) | 968
- [Options](#) | 968
- [Required Privilege Level](#) | 969
- [Release Information](#) | 969

### Syntax

```
link-protection {
 non-revertive;
}
```

### Hierarchy Level

```
[edit chassis aggregated-devices ethernet lacp]
```

### Description

Enable LACP link protection at the global (chassis) level.

### Options

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.3.

## RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices](#)

## link-speed (Aggregated Ethernet)

### IN THIS SECTION

- [Syntax | 969](#)
- [Hierarchy Level \(EX Series\) | 970](#)
- [Description | 970](#)
- [Options | 970](#)
- [Required Privilege Level | 971](#)
- [Release Information | 971](#)

## Syntax

```
link-speed speed;
```

## Hierarchy Level (EX Series)

```
[edit interfaces aex aggregated-ether-options],
[edit interfaces interface-range name aggregated-ether-options],
[edit interfaces interface-range name aggregated-sonet-options]
```

## Description

For aggregated Ethernet interfaces only, set the required link speed.

## Options

**speed**—For aggregated Ethernet links, you can specify **speed** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.

Aggregated Ethernet links on T Series, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.

- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Enables bundling of different Ethernet rate links in the same Aggregated Ethernet interface.
- **oc192**—Links are OC192.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### Release Information

Statement introduced before Junos OS Release 7.4.

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

### RELATED DOCUMENTATION

---

[Aggregated Ethernet Interfaces Overview](#)

---

[Configuring Aggregated Ethernet Link Speed](#)

---

[Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles](#)

---

[Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)

---

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

## link-speed (Aggregated SONET/SDH)

### IN THIS SECTION

- [Syntax | 972](#)
- [Hierarchy Level | 972](#)
- [Description | 972](#)
- [Options | 972](#)
- [Required Privilege Level | 973](#)
- [Release Information | 973](#)

### Syntax

```
link-speed (speed | mixed);
```

### Hierarchy Level

```
[edit interfaces asx aggregated-sonet-options]
```

### Description

For aggregated SONET/SDH interfaces only, set the required link speed.

### Options

*speed*—Aggregated SONET/SDH links can have one of the following speed values.

- **oc3**—Links are OC3c or STM1c.
- **oc12**—Links are OC12c or STM4c.
- **oc48**—Links are OC48c or STM16c.
- **oc192**—Links are OC192c or STM64c.

- **oc768**—Links are OC768c or STM256c.

**mixed**—For aggregated SONET/SDH links on T Series routers, you can mix interface speeds in SONET/SDH aggregation bundles. Interface speeds from OC3 through OC768 are supported.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### Release Information

Statement introduced before Junos OS Release 7.4.

**mixed** option added in Release 8.0.

### RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Link Speed](#)

[Configuring Aggregated SONET/SDH Interfaces](#)

## Imi (Ethernet OAM)

### IN THIS SECTION

- [Syntax | 974](#)
- [Hierarchy Level | 974](#)
- [Description | 974](#)
- [Options | 974](#)
- [Required Privilege Level | 975](#)
- [Release Information | 975](#)

## Syntax

```
lmi {
 status-counter count;
 polling-verification-timer value;
 interface name {
 uni-id uni-name;
 status-counter number;
 polling-verification-timer value;
 evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
 evc evc-name {
 default-evc;
 vlan-list vlan-id-list;
 }
 }
}
```

## Hierarchy Level

```
[edit protocols oam ethernet]
```

## Description

On routers with **ge**, **xe**, or **ae** interfaces, configure an OAM Ethernet Local Management Interface (E-LMI).

**NOTE:** On MX Series routers, E-LMI is supported on Gigabit Ethernet (**ge**), 10-Gigabit Ethernet (**xe**), and Aggregated Ethernet (**ae**) interfaces configured on MX Series routers with DPC only.

## Options

**status-counter** *count*—Status counter (N393), defaults to 4.

**interface** *name*—Polling verification timer (T392), defaults to 15 seconds.

**uni-id** *uni-name*—(Optional) Defaults to the physical interface name.

**status-counter** *number*—(Optional) Defaults to a global value.

**polling-verification-timer *value***—(Optional) Defaults to a global value.

**evc-map-type (all-to-one-bundling | bundling | service-multiplexing)**—Specify the Ethernet virtual connection (EVC) map type.

**evc *evc-name***—Specify the name of the EVC.

**default-evc**—Set the specified EVC as the default EVC.

**vlan-list *vlan-id-list***—Specify a group of VLANs to assign to the EVC.

## 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 Release 9.5.

## RELATED DOCUMENTATION

*Configuring Ethernet Local Management Interface*

[evcs](#) | [869](#)

## load-balance

### IN THIS SECTION

- [Syntax](#) | [976](#)
- [Hierarchy Level](#) | [976](#)
- [Description](#) | [976](#)
- [Options](#) | [977](#)
- [Required Privilege Level](#) | [977](#)
- [Release Information](#) | [977](#)

## Syntax

```
load-balance {
 adaptive{
 pps;
 scan-interval multiple;
 tolerance percentage;
 }
 no-adaptive;
 per-packet;
}
```

## Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options],
[edit dynamic-profiles name interfaces name logical-tunnel-options],
[edit dynamic-profiles name interfaces interface-range name aggregated-ether-
options],
[edit dynamic-profiles name interfaces interface-range name logical-tunnel-
options],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-
ether-options],
[edit dynamic-profiles name logical-systems name interfaces name logical-tunnel-
options],
[edit dynamic-profiles name logical-systems name interfaces interface-range name
aggregated-ether-options],
[edit dynamic-profiles name logical-systems name interfaces interface-range name
logical-tunnel-options],
[edit interfaces name aggregated-ether-options],
[edit interfaces name logical-tunnel-options],
[edit interfaces interface-range name aggregated-ether-options],
[edit interfaces interface-range name logical-tunnel-options]
```

## Description

Load-balances the received traffic across all the available paths of aggregated Ethernet bundles for better link utilization.

## Options

- adaptive** (MX Series and PTX Series) Corrects a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an Aggregated Ethernet bundle.
- no-adaptive** (MX Series and PTX Series) Disables the adaptive load-balancing solution configured to distribute traffic by using a feedback mechanism.
- per-packet** (MX Series only) Randomly sprays packets to the aggregate next hops in a round-robin manner to avoid traffic imbalance.

## Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 13.3.

## RELATED DOCUMENTATION

*Understanding Aggregated Ethernet Load Balancing*

*Example: Configuring Aggregated Ethernet Load Balancing*

## load-balance-stateful (Aggregated Ethernet Interfaces)

### IN THIS SECTION

- [Syntax | 978](#)
- [Hierarchy Level | 978](#)
- [Description | 978](#)
- [Options | 978](#)
- [Required Privilege Level | 978](#)

## Syntax

```
load-balance-stateful {
 per-flow;
 rebalance interval;
 load-type (low | medium | large);
}
```

## Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options]
```

## Description

Define 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.

## Options

**stateful** Define the stateful load-distribution mechanism for traffic flows on aggregated Ethernet interfaces.

## 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 Release 13.2R1.

## RELATED DOCUMENTATION

| *Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces*

## load-type (Aggregated Ethernet Interfaces)

### IN THIS SECTION

- [Syntax | 979](#)
- [Hierarchy Level | 979](#)
- [Description | 980](#)
- [Options | 980](#)
- [Required Privilege Level | 980](#)
- [Release Information | 980](#)

## Syntax

```
load-type (low | medium | large);
```

## Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-
stateful]
```

## Description

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.

## Options

- low** Define a low load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is less or minimal (between 1 and 100 flows).
- medium** Define a medium or moderate load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is relatively higher (between 100 and 1000 flows).
- large** Define a high load-balancing method if the number of flows that flow on the specified aggregated Ethernet interface is excessive or reaches the maximum supported flows (between 1000 and 10,000 flows).

## 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 Release 13.2R1.

## RELATED DOCUMENTATION

| *Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces*

## local-bias (ae load-balance)

### IN THIS SECTION

- [Syntax | 981](#)
- [Hierarchy Level | 981](#)
- [Description | 981](#)
- [Required Privilege Level | 982](#)
- [Release Information | 982](#)

### Syntax

```
local-bias percent bias;
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options load-balance]
```

### Description

Next hop addresses may be local or remote, and traffic can be expected to be more-or-less evenly distributed among the available next-hop addresses whether they are local or remote. You can skew distribution to favor local addresses, however, by setting a value for local bias (local relative to the packet forwarding engine (PFE) performing the packet look up).

For example, a value of 100 would exclude remote next-hop addresses from the traffic distribution by forcing 100% of next-hop traffic flows to use local addresses. A value of 50, on the other hand, would skew 50% of the flows that would otherwise use remote links so they use local links instead. That is, for a value set to 50, given four next-hop links, two of which are local and two of which are remote, each of the remote links could be expected to get one eighth of the flows ( $25\% / 2$ ) = 12.5%. Likewise, each of the local links could also be expected to receive about a third of the flows ( $25\% + 12.5\%$ ) = 37.5%.

In contrast, with no value set for local bias, each of the four links would be expected to receive 25% of the total flows.

You can use **local-bias** with adaptive load balancing, which uses a feedback mechanism to automatically correct load imbalance by adjusting the bandwidth and packet streams traversing links within an AE bundle. In this case, **local-bias** also employs a combination of link-saturation prediction and random restart delay to govern link utilization in a way that prevents oscillation of the load and load balancing schemes in effect.

**NOTE:** MPC5 and MPC6 line cards include XM and XL-based packet forwarding engines, or PFEs, and locality is decided on the basis of the XL chip, not the XM chip. Therefore, when an AE bundle has child links hosted on two different XMs that are connected (in the chip architecture) to a single XL, they are considered local to the XL PFE. In practice, what this means is that if a single AE interface includes member links that happen to be spread over two XMs but are actually served by the same XL, **local-bias** may not work as expected because links are considered local to the XL PFE.

## Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 19.2R1.

## RELATED DOCUMENTATION

*Understanding Aggregated Ethernet Load Balancing*

*Example: Configuring Aggregated Ethernet Load Balancing*

## logical-tunnel-options

### IN THIS SECTION

● [Syntax | 983](#)

● [Hierarchy Level | 983](#)

- Description | 983
- Options | 984
- Required Privilege Level | 984
- Release Information | 984

## Syntax

```

logical-tunnel-options {
 link-protection {
 non-revertive;
 revertive;
 }
 load-balance{
 adaptive{
 pps;
 scan-interval scan-interval;
 tolerance percent;
 }
 no-adaptive;
 per-packet;
 local-biaspercent;
 }
 per-unit-mac-disable;
}

```

## Hierarchy Level

```

[edit dynamic-profiles name interfaces],
[edit dynamic-profiles name logical-systems name interfaces],
[edit interfaces]

```

## Description

For redundant logical tunnels, specifies the logical tunnel interface-specific options for load balancing and link protection. The remaining statements are explained separately. See [CLI Explorer](#).

## Options

**link-protection** Enables link protection for redundant logical tunnel interfaces. In addition to enabling static link protection, you must configure a primary and secondary (backup) link for egress traffic.

- Values:
  - non-revertive—Do not revert back from active backup link to primary, if primary is UP.
  - revertive—Revert back from active backup link to primary, if primary is UP.
- **Default:** revertive

**per-unit-mac-disable** Disable the creation of per unit mac address on LT IFLs for VPLS/CCC encaps

## Required Privilege Level

interface

## Release Information

Statement introduced in Junos OS Release 19.2R1.

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

### IN THIS SECTION

- Syntax | 985
- Hierarchy Level | 985
- Description | 985
- Default | 986
- Required Privilege Level | 986
- Release Information | 986

## Syntax

```
(loopback | no-loopback);
```

## Hierarchy Level

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options],
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigheter-options],
[edit interfaces interface-range name ether-options]
```

For QFX Series and EX Series:

```
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options],
```

For SRX Series Devices and vSRX:

```
[edit interfaces interface-name redundant-ether-options]
```

## Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.

### NOTE:

- By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.
- IPv6 Neighbor Discovery Protocol (NDP) addresses are not supported on Gigabit Ethernet interfaces when loopback mode is enabled on the interface. That is, if the **loopback** statement is configured at the `[edit interfaces ge-fpc/pic/port gigheter-options]` hierarchy level, an NDP address cannot be configured at the `[edit interfaces ge-fpc/pic/port unit logical-unit-number family inet6 address]` hierarchy level.

## Default

By default, loopback is disabled.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement modified in Junos OS Release 9.2 for the SRX Series.

## RELATED DOCUMENTATION

[Configuring Ethernet Loopback Capability](#)

[Understanding Interfaces](#)

## loopback (Local and Remote)

### IN THIS SECTION

- [Syntax | 987](#)
- [Hierarchy Level | 987](#)
- [Description | 987](#)
- [Options | 987](#)
- [Required Privilege Level | 987](#)
- [Release Information | 987](#)

## Syntax

```
loopback (local | remote);
```

## Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

## Description

Enables local loopback and enables remote loopback. This allows you to test the transceiver cable connection from the far end to the retimer interface without changing the cable.

## Options

|               |                         |
|---------------|-------------------------|
| <b>local</b>  | Enables local loopback  |
| <b>remote</b> | Enables remote loopback |

## 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 Release 15.1F3 and 16.1R2.

## RELATED DOCUMENTATION

| [Configuring Ethernet Loopback Capability](#)

## loopback-tracking

### IN THIS SECTION

- [Syntax | 988](#)
- [Hierarchy Level | 988](#)
- [Description | 988](#)
- [Required Privilege Level | 988](#)
- [Release Information | 988](#)

### Syntax

```
loopback-tracking;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

### Description

Enables loopback tracking on Ethernet interfaces. When loopback tracking is enabled and the Ethernet Operation, Administration, and Management (OAM) link-fault management process (lfmd) detects its own generated packets on an interface, it marks the interface as down. When the loopback issue resolves, the interface is brought back up.

### 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 Release 14.2.

## RELATED DOCUMENTATION

[IEEE 802.3ah OAM Link-Fault Management Overview](#)

[Enabling IEEE 802.3ah OAM Support](#)

## loss-priority

### IN THIS SECTION

- [Syntax | 989](#)
- [Hierarchy Level | 989](#)
- [Description | 989](#)
- [Options | 990](#)
- [Required Privilege Level | 990](#)
- [Release Information | 990](#)

### Syntax

```
loss-priority (high | low);
```

### Hierarchy Level

```
[edit interfaces interface-name together-options ethernet-switch-profile
ethernet-policer-profile output-priority-map classifier premium forwarding-class
class-name]
```

### Description

Specify the packet loss priority value.

## Options

**high**—Packet has high loss priority.

**low**—Packet has low loss priority.

## Required Privilege Level

**interface**—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Configuring Gigabit Ethernet Policers](#) | 271

## loopback-remote

### IN THIS SECTION

- [Syntax](#) | 990
- [Hierarchy Level](#) | 991
- [Description](#) | 991
- [Required Privilege Level](#) | 991
- [Release Information](#) | 991

## Syntax

```
loopback-remote;
```

## Hierarchy Level

```
[edit dynamic-profiles name interfaces name gigheter-options],
[edit dynamic-profiles name logical-systems name interfaces name gigheter-
options],
[edit interfaces name gigheter-options]
```

## Description

Configure **loopback-remote** on the remote end of the link. When you configure on an interface **loopback-remote**, the local end goes down and remote end goes up.

## Required Privilege Level

interface

## Release Information

### RELATED DOCUMENTATION

| [Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#) | [681](#)

## mac

### IN THIS SECTION

- [Syntax](#) | [992](#)
- [Hierarchy Level](#) | [992](#)
- [Description](#) | [992](#)
- [Options](#) | [992](#)
- [Required Privilege Level](#) | [992](#)
- [Release Information](#) | [992](#)

## Syntax

```
mac mac-address;
```

## Hierarchy Level

```
[edit interfaces interface-name]
```

## Description

Set the MAC address of the interface.

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

## Options

***mac-address***—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nr.nr.nr.nr.nr.nr*. For example, 0000.5e00.5355 or 00:00:5e:00:53:55.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring the MAC Address on the Management Ethernet Interface](#)

[Configuring a Pseudowire Subscriber Logical Interface Device](#)

## mac-address (Accept Source Mac)

### IN THIS SECTION

- [Syntax | 993](#)
- [Hierarchy Level | 993](#)
- [Description | 993](#)
- [Options | 993](#)
- [Required Privilege Level | 994](#)
- [Release Information | 994](#)

### Syntax

```
mac-address mac-address policer;
```

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number accept-source-mac],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number accept-source-mac]
```

### Description

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, and 100-Gigabit Ethernet Type 5 PIC with CFP, specify a remote MAC address on which to count incoming and outgoing packets.

### Options

***mac-address***—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn.nn.nn.nn.nn.nn*. For example, 0011.2233.4455 or 00:11:22:33:44:55.

***policer***—MAC policer. For more information, see ["policer \(MAC\)" on page 1076](#).

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

## RELATED DOCUMENTATION

| [Configuring Gigabit Ethernet Policers](#) | 271

## mac-learn-enable

### IN THIS SECTION

- [Syntax](#) | 994
- [Hierarchy Level](#) | 995
- [Description](#) | 995
- [Required Privilege Level](#) | 995
- [Release Information](#) | 995

## Syntax

```
mac-learn-enable;
```

## Hierarchy Level

```
[edit interfaces interface-name together-options ethernet-switch-profile]
[edit interfaces aex aggregated-ether-options ethernet-switch-profile]
```

## Description

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, configure dynamic learning of the source and destination MAC addresses. By default, the interface is not allowed to dynamically learn source and destination MAC addresses.

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

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Support for statement under the **[edit interfaces aex aggregated-ether-options ethernet-switch-profile]** hierarchy introduced in Junos OS Release 15.1.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[Configuring MAC Address Accounting | 23](#)

## mac-validate

### IN THIS SECTION

- [Syntax | 996](#)
- [Hierarchy Level | 996](#)
- [Description | 996](#)
- [Options | 996](#)
- [Required Privilege Level | 997](#)
- [Release Information | 997](#)

### Syntax

```
mac-validate (loose | strict);
```

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family family]
```

### Description

Enable IP and MAC address validation for static Ethernet and IP demux interfaces.

### Options

**loose**—Forwards incoming 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 match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.

**strict**—Forwards incoming 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.

## 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 Release 9.3.

## RELATED DOCUMENTATION

[MAC Address Validation on Static Ethernet Interfaces Overview](#)

[Configuring an IP Demultiplexing Interface](#)

[Configuring a VLAN Demultiplexing Interface](#)

## master-only

### IN THIS SECTION

- [Syntax | 997](#)
- [Hierarchy Level | 998](#)
- [Description | 998](#)
- [Required Privilege Level | 998](#)
- [Release Information | 998](#)

## Syntax

```
master-only;
```

## Hierarchy Level

```
[edit groups rex interfaces (fxp0 | em0) unit logical-unit-number family family
address],
[edit groups rex logical-systems logical-system-name interfaces fxp0 unit
logical-unit-number family family address],
[edit interfaces (fxp0 | em0) unit logical-unit-number family family address],
[edit logical-systems logical-system-name interfaces fxp0 unit logical-unit-
number family family address]
```

## Description

Configure the IP address to be used when the Routing Engine is the current primary.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring a Consistent Management IP Address | 26](#)

[Junos OS CLI User Guide](#)

## max-sessions (PPPoE Service Name Tables)

### IN THIS SECTION

● [Syntax | 999](#)

● [Hierarchy Level | 999](#)

- Description | 999
- Options | 999
- Required Privilege Level | 999
- Release Information | 1000

## Syntax

```
max-sessions number;
```

## Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name]
```

## Description

Configure the maximum number of active PPPoE sessions using either static or dynamic PPPoE interfaces that the router can establish with the specified named service, **empty** service, or **any** service entry in a PPPoE service name table. The router maintains a count of active PPPoE sessions for each service entry to determine when the maximum sessions limit has been reached.

The router uses the **max-sessions** value for a PPPoE service name table entry in conjunction with the **max-sessions** value configured for the PPPoE underlying interface, and with the maximum number of PPPoE sessions supported on your router. If your configuration exceeds any of these maximum session limits, the router is unable to establish the PPPoE session.

## Options

***number***—Maximum number of active PPPoE sessions that the router can establish with the specified PPPoE service name table entry, in the range 1 to the platform-specific maximum PPPoE sessions supported for your router. The default value is equal to the maximum number of PPPoE sessions supported on your routing platform.

## 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 Release 10.2.

## RELATED DOCUMENTATION

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

*Configuring PPPoE Service Name Tables*

*PPPoE Maximum Session Limit Overview*

*Configuring an Interface Set of Subscribers in a Dynamic Profile*

*Subscriber Interfaces and PPPoE Overview*

## max-sessions-vsa-ignore (Static and Dynamic Subscribers)

### IN THIS SECTION

- [Syntax | 1000](#)
- [Hierarchy Level | 1001](#)
- [Description | 1001](#)
- [Default | 1001](#)
- [Required Privilege Level | 1001](#)
- [Release Information | 1001](#)

## Syntax

```
max-sessions-vsa-ignore;
```

## Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number
family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-
number family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-
number pppoe-underlying-options],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-
options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number pppoe-underlying-options]
```

## Description

Configure the router to ignore (clear) the value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks vendor-specific attribute (VSA) [26-143], and restore the PPPoE maximum session value on the underlying interface to the value configured in the CLI with the **max-sessions** statement. The PPPoE maximum session value specifies the maximum number of concurrent static or dynamic PPPoE logical interfaces (sessions) that the router can activate on the PPPoE underlying interface, or the maximum number of active static or dynamic PPPoE sessions that the router can establish with a particular service entry in a PPPoE service name table.

## Default

If you do not include the **max-sessions-vs-a-ignore** statement, the maximum session value returned by RADIUS in the Max-Clients-Per-Interface VSA takes precedence over the PPPoE maximum session value configured with the **max-sessions** statement.

## 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 Release 11.4.

## RELATED DOCUMENTATION

*Limiting the Maximum Number of PPPoE Sessions on the Underlying Interface*

*PPPoE Maximum Session Limit Overview*

*Guidelines for Using PPPoE Maximum Session Limit from RADIUS*

*Juniper Networks VSAs Supported by the AAA Service Framework*

*Configuring an Interface Set of Subscribers in a Dynamic Profile*

*Subscriber Interfaces and PPPoE Overview*

## maximum-links

### IN THIS SECTION

- [Syntax | 1002](#)
- [Hierarchy Level | 1002](#)
- [Description | 1002](#)
- [Options | 1003](#)
- [Required Privilege Level | 1003](#)
- [Release Information | 1003](#)

### Syntax

```
maximum-links maximum-links-limit;
```

### Hierarchy Level

```
[edit chassis aggregated-devices]
```

### Description

Configure the maximum links limit for aggregated devices. Note that for MX Series routers, to set a range of 32 or 64 the router must be running in Enhanced IP mode, which is only supported for Trio-

based MPCs and multiservice DPCs (MS-DPCs). For more information on Enhanced IP mode, [Network Services Mode Overview](#).

For MX series routers and PTX series switches, the option for 64 links is only supported for Junos OS release 12.3 and later.

**NOTE:** This statement is not supported on the MX80, MX104, and PTX1000 routers.

## Options

*maximum-links-limit* Maximum links limit for aggregated devices.

- **Range:** 16, 32, 64

**NOTE:** On T-Series routers, the maximum-links supported is 32 in an aggregated Ethernet link.

## 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 Release 11.1.

## RELATED DOCUMENTATION

[Network Services Mode Overview](#)

[Configuring Junos OS for Supporting Aggregated Devices](#)

[Configuring an Aggregated Ethernet Interface](#)

*network-services*

## mc-ae

### IN THIS SECTION

- [Syntax | 1004](#)
- [Hierarchy Level | 1004](#)
- [Description | 1005](#)
- [Options | 1005](#)
- [Required Privilege Level | 1009](#)
- [Release Information | 1009](#)

### Syntax

```
mc-ae {
 chassis-id chassis-id;
 events {
 iccp-peer-down;
 force-icl-down;
 prefer-status-control-active;
 }
 init-delay-time seconds;
 mc-ae-id mc-ae-id;
 mode (active-active | active-standby);
 redundancy-group group-id;
 revert-time revert-time;
 status-control (active | standby);
 switchover-mode (non-revertive | revertive);
}
```

### Hierarchy Level

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

## Description

Enable multichassis link aggregation groups (MC-LAG), which enables one device to form a logical LAG interface with two or more other devices.

## Options

|                                     |                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>chassis-id</b>                   | Specify the chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links. Each MC-LAG peer should have a unique chassis ID. <ul style="list-style-type: none"> <li>• <b>Values:</b> 0 or 1</li> </ul>                                                                                                                                            |
| <b>events</b>                       | Specify an action if a specific MC-LAG event occurs.                                                                                                                                                                                                                                                                                                                                                     |
| <b>iccp-peer-down</b>               | Specify an action if the ICCP peer of this node goes down.                                                                                                                                                                                                                                                                                                                                               |
| <b>force-icl-down</b>               | If the node's ICCP peer goes down, bring down the interchassis-link logical interface.                                                                                                                                                                                                                                                                                                                   |
| <b>prefer-status-control-active</b> | Specify that the node configured as <b>status-control active</b> become the active node if the peer of this node goes down.<br><br>When ICCP goes down, you can use this keyword to make a mc-lag PE to become the active PE. For example, if you want mc-lag PE1 to be Active on ICCP down, then configure this keyword in PE1. It is not recommended to configure this keyword in both the mc-lag PEs. |

**NOTE:** The **prefer-status-control-active** statement can be configured with the **status-control standby** configuration to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP will not go down unless the router or switch is down. You must also configure the **hold-time down** value (at the **[edit interfaces *interface-name*]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents data traffic loss by ensuring that when the router or switch with the **status-control active** configuration goes down, the router or switch with the **status-control standby** configuration does not go into standby mode.

To make the **prefer-status-control-active** configuration work with the **status-control standby** configuration when an interchassis-link logical interface is configured on aggregate Ethernet interface, you must either configure the **lACP periodic interval** statement at the **[edit interface *interface-name* aggregated-ether-options]** hierarchy level as **slow** or configure the **detection-time threshold** statement at the **[edit protocols iccp peer liveness-detection]** hierarchy level as less than 3 seconds.

**init-delay-time**  
*seconds* To minimize traffic loss, specify the number of seconds in which to delay bringing the multichassis aggregated Ethernet interface back to the up state when you reboot an MC-LAG peer. By delaying the startup of the interface until after protocol convergence, you can prevent packet loss during the recovery of failed links and devices.

**NOTE:** On QFX and EX Series switches, the default session establishment hold time is 300 seconds. However, the session establishment time must be at least 100 seconds higher than the init delay time. You can optionally update the session establishment time to be 340 seconds and the init delay time to be 240 seconds.

**mc-ae-id** *mc-ae-id* Specify the identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number.

- **Range:** 1 through 65,535

**mode** (**active-active** | **active-standby**) Specify whether the MC-LAG is in active-active or active-standby mode. Chassis that are in the same group must be in the same mode.

**NOTE:** You can configure IPv4 (*inet*) and IPv6 (*inet6*) addresses on **mc-ae** interfaces when the **active-standby** mode is configured.

In active-active mode, all member links are active on the MC-LAG. In this mode, media access control (MAC) addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer. Active-active mode is a simple and deterministic design and is easier to troubleshoot than active-standby mode.

**NOTE:** Active-active mode is not supported on Dense Port Concentrator (DPC) line cards. Instead, use active-standby mode.

In active-active MC-LAG topologies, network interfaces are categorized into three interface types, as follows:

- S-Link—Single-homed link (S-Link) terminating on an MC-LAG peer device
- MC-Link—MC-LAG link
- ICL—Inter-chassis link
- Mode

Indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

In active-active mode, all member links are active on the MC-LAG. In this mode, media access control (MAC) addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer. Active-active mode is a simple and deterministic design and is easier to troubleshoot than active-standby mode.

**NOTE:** Active-active mode is not supported on Dense Port Concentrator (DPC) line cards. Instead, use active-standby mode.

Depending on the incoming and outgoing interface types, some constraints are added to the Layer 2 forwarding rules for MC-LAG configurations. The following data traffic forwarding rules apply.

**NOTE:** If only one MC-LAG member link is in the up state, it is considered an S-Link.

- When an MC-LAG network receives a packet from a local MC-Link or S-Link, the packet is forwarded to other local interfaces, including S-Links and MC-Links based on the normal Layer 2 forwarding rules and on the configuration of the **mesh-group** and **no-local-switching** statements. If MC-Links and S-Links are

in the same mesh group and their **no-local-switching** statements are enabled, the received packets are only forwarded upstream and not sent to MC-Links and S-Links.

- The following circumstances determine whether or not an ICL receives a packet from a local MC-Link or S-Link:
  - If the peer MC-LAG network device has S-Links or MC-LAGs that do not reside on the local MC-LAG network device
  - Whether or not interfaces on two peering MC-LAG network devices are allowed to talk to each other
- When an MC-LAG network receives a packet from the ICL, the packet is forwarded to all local S-Links and active MC-LAGs that do not exist in the MC-LAG network from which the packet was sent.

In active-standby mode, only one of the MC-LAG peers is active at any given time. The other MC-LAG peer is in backup (standby) mode. The active MC-LAG peer uses Link Aggregation Control Protocol (LACP) to advertise to client devices that its child link is available for forwarding data traffic. Active-standby mode should be used if you are interested in redundancy only. If you require both redundancy and load sharing across member links, use active-active mode.

**NOTE:** Active-standby mode is not supported on EX4300 and QFX Series switches.

**redundancy-  
group *group-id***

Specify the redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions. The value you assign to this option must match the value you assign to the **redundancy-group-id-list** option that you configure on the ICCP peer. If the value differs, when you commit the changes, the system reports a commit check error.

**BEST PRACTICE:** We recommend that you configure only one redundancy group between MC-LAG nodes. The redundancy group represents the domain of high availability between the MC-LAG nodes. One redundancy group is sufficient between a pair of MC-LAG nodes. If you are using logical systems, then configure one redundancy group between MC-LAG nodes in each logical system.

|                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                    | <ul style="list-style-type: none"> <li>• <b>Range:</b> 1 through 4,294,967,294</li> </ul>                                                                                                                                                                                                                                                                                                                                                                         |
| <b>revert-time</b>                                                 | <p>Wait interval (in minutes) before the switchover to the preferred node is performed when the <b>switchover-mode</b> is configured as revertive.</p> <ul style="list-style-type: none"> <li>• <b>Range:</b> 1 through 10</li> </ul>                                                                                                                                                                                                                             |
| <b>status-control</b><br>( <b>active</b>   <b>standby</b> )        | <p>Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.</p> <ul style="list-style-type: none"> <li>• Events ICCP-Peer-Down Force-ICL-Down<br/>Forces the ICL to be down if the peer of this node goes down.</li> <li>• Events ICCP-Peer-Down Prefer-Status-Control-Active<br/>Allows the LACP system ID to be retained during a reboot, which provides better convergence after a failover.</li> </ul> |
| <b>switchover-mode</b> ( <b>non-revertive</b>   <b>revertive</b> ) | <p>Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.</p>                                                                                                                                                                                                                                                                                                                                         |

**NOTE:** For **revertive** mode to automatically switch over to the preferred node, the **status-control** statement should be configured as **active**.

## 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 Release 9.6.

**events** statement introduced in Junos OS Release 11.4R4 for MX Series routers.

**prefer-status-control-active** statement introduced in Junos OS Release 13.2R1 for EX Series switches.

**init-delay-time** *seconds* statement introduced in Junos OS Release 13.2R3 for EX Series switches.

**switchover-mode** and **revert-time** statements introduced in Junos OS Release 13.3.

Support for logical systems introduced in Junos OS Release 14.1.

## minimum-bandwidth (aggregated Ethernet)

### IN THIS SECTION

- Syntax | 1010
- Hierarchy Level | 1010
- Description | 1010
- Options | 1011
- Required Privilege Level | 1011
- Release Information | 1011

### Syntax

```
minimum-bandwidth bw-unit unit bw-value value;
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

### Description

Configure the minimum bandwidth unit for an aggregated Ethernet bundle as bps, Gbps, Kbps, or Mbps and the bandwidth value from 1 through 128,000.

(T Series, MX Series, PTX Series routers) You cannot configure the minimum number of links and the minimum bandwidth for an aggregated Ethernet bundle at the same time. They are mutually exclusive. To determine the status of the bundle, the device compares the value configured for minimum links and the value for minimum bandwidth. Because both cannot be configured at the same time, the device compares the configured value of the parameter with the default value of the other parameter. The device picks the higher value of the two parameters to determine the status of the Ethernet bundle. [Table 130 on page 1011](#) describes how the device determines the bundle status based on sample values assigned to both parameters

**Table 130: Determination of Bundle Status based on Minimum links and Minimum bandwidth Parameters**

| Current Bandwidth | Minimum Links                           | Minimum Bandwidth      | Bundle Status |
|-------------------|-----------------------------------------|------------------------|---------------|
| 100G<br>(10x10G)  | 2<br>Bandwidth=20G (2x10G)              | N/A<br>Default (1 bps) | Up            |
| 50G<br>(5x10G)    | N/A<br>Default (1 link) Bandwidth = 10G | 20G                    | Up            |
| 50G<br>(5x10G)    | N/A<br>Default (1 link) Bandwidth = 10G | 60G                    | Down          |

## Options

*unit*—Minimum bandwidth unit for the aggregated Ethernet bundle as bps, Gbps, Kbps, or Mbps.

*value*—Minimum bandwidth value from 1 through 128,000.

- **Default:** 1

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 14.1R1 and 14.2.

## RELATED DOCUMENTATION

[Aggregated Ethernet Interfaces Overview](#)

[Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles](#)

[Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles](#)

## minimum-links

### IN THIS SECTION

- [Syntax \(SRX, MX, T, M, EX, QFX Series, EX4600, Qfabric System\) | 1012](#)
- [Hierarchy Level \(EX Series\) | 1012](#)
- [Hierarchy Level \(QFX Series\) | 1012](#)
- [Description | 1013](#)
- [Options | 1013](#)
- [Required Privilege Level | 1014](#)
- [Release Information | 1014](#)

### Syntax (SRX, MX, T, M, EX, QFX Series, EX4600, Qfabric System)

```
minimum-links number;
```

### Hierarchy Level (EX Series)

```
[edit interfaces aex aggregated-ether-options],
[edit interfaces aex aggregated-sonet-options],
[edit interfaces interface-name mlfr-uni-nni-bundle-options],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range range aggregated-ether-options],
[edit interfaces interface-range range aggregated-sonet-options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

### Hierarchy Level (QFX Series)

```
[edit interfaces aex aggregated-ether-options]
```

## Description

For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.

(T Series, MX Series, PTX Series routers) You cannot configure the minimum number of links and the minimum bandwidth of an aggregated Ethernet bundle at the same time. They are mutually exclusive. To determine the status of the bundle, the device compares the value configured for minimum links and the value for minimum bandwidth. Because both cannot be configured at the same time, the device compares the configured value of the parameter with the default value of the other parameter. The device picks the higher value of the two parameters to determine the status of the Ethernet bundle.

[Table 131 on page 1013](#) describes how the device determines the bundle status based on sample values assigned to both parameters

**Table 131: Determination of Bundle Status based on Minimum links and Minimum bandwidth Parameters**

| Current Bandwidth | Minimum Links                           | Minimum Bandwidth      | Bundle Status |
|-------------------|-----------------------------------------|------------------------|---------------|
| 100G<br>(10x10G)  | 2<br>Bandwidth=20G (2x10G)              | N/A<br>Default (1 bps) | Up            |
| 50G<br>(5x10G)    | N/A<br>Default (1 link) Bandwidth = 10G | 20G                    | Up            |
| 50G<br>(5x10G)    | N/A<br>Default (1 link) Bandwidth = 10G | 60G                    | Down          |

## Options

*number*—Number of links.

- **Range:** On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. 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 EX4600, QFX Series and Q Fabric Systems , 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.

- **Default:** 1

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Minimum Links | 109](#)

[Configuring Aggregated SONET/SDH Interfaces](#)

[Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)

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

[Junos OS Services Interfaces Library for Routing Devices](#)

*Configuring Link Aggregation*

## mixed-rate-mode

### IN THIS SECTION

● [Syntax | 1015](#)

● [Hierarchy Level | 1015](#)

- [Description | 1015](#)
- [Required Privilege Level | 1015](#)
- [Release Information | 1015](#)

## Syntax

mixed-rate-mode;

## Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number mixed-rate-mode],
[edit chassis lcc number fpc slot-number pic pic-number mixed-rate-mode]
(Routing Matrix)
```

## Description

Configure the mixed-rate mode for the 24-port 10 Gigabit Ethernet PIC (PF-24XGE-SFPP) only.

## 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 Release 13.3.

## RELATED DOCUMENTATION

[Modes of Operation | 204](#)

[Configuring Mixed-Rate Mode Operation | 205](#)

## mtu

### IN THIS SECTION

- [Syntax | 1016](#)
- [Hierarchy Level | 1016](#)
- [Description | 1017](#)
- [Options | 1019](#)
- [Required Privilege Level | 1019](#)
- [Release Information | 1019](#)

### Syntax

```
mtu bytes;
```

### Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number family family],
[edit interfaces interface-range name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family],
[edit logical-systems logical-system-name protocols l2circuit local-switching interface interface-name backup-neighbor address],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface interface-name],
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface interface-name backup-neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls],
[edit protocols l2circuit local-switching interface interface-name backup-neighbor address],
[edit protocols l2circuit neighbor address interface interface-name]
```

```
[edit protocols l2circuit neighbor address interface interface-name backup-
neighbor address],
[edit routing-instances routing-instance-name protocols l2vpn interface
interface-name],
[edit routing-instances routing-instance-name protocols vpls],
[edit logical-systems name protocols ospf area name
interface],
[edit logical-systems name routing-instances
name protocols ospf area name interface],
[edit protocols ospf area name interface],
[edit routing-instances name protocols ospf area name interface]
```

## Description

Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

To route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named `irb` or `vlan`, respectively).



**CAUTION:** For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.

**NOTE:** The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the **flexible-vlan-tagging** statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
- In case of Layer 2 IFL configured with the **vlan-tagging** statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.

**NOTE:**

- If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.
- Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.
- On ACX Series routers, you can configure the protocol MTU by including the **mtu** statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] or [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level.
  - If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
  - If you are configuring the protocol MTU for both **inet** and **inet6** families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for **inet** and **inet6** families that are configured on the same logical interface.
- Starting in Release 14.2, MTU for IRB interfaces is calculated by removing the Ethernet header overhead ( $6(\text{DMAC})+6(\text{SMAC})+2(\text{EtherType})$ ), and the MTU is a minimum of the two values:
  - Configured MTU
  - Associated bridge domain's physical or logical interface MTU
    - For Layer 2 logical interfaces configured with **flexible-vlan-tagging**, IRB MTU is calculated by including 8 bytes overhead ( $\text{SVLAN}+\text{CVLAN}$ ).
    - For Layer 2 logical interfaces configured with **vlan-tagging**, IRB MTU is calculated by including single VLAN 4 bytes overhead.

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

For more information about configuring MTU for specific interfaces and router or switch combinations, see [Configuring the Media MTU](#).

## Options

*bytes*—MTU size.

- **Range:** 256 through 9192 bytes, 256 through 9216 (EX Series switch interfaces), 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers), 256 through 9500 bytes (Junos OS 16.1R1 for MX Series routers)

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

- MPC1
  - MPC2
  - MPC2E
  - MPC3E
  - MPC4E
  - MPC5E
  - MPC6E
- **Default:** 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Support at the[set interfaces interface-name unit logical-unit-number family ccc] hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.

## RELATED DOCUMENTATION

[Configuring the Media MTU](#)

*Configuring the MTU for Layer 2 Interfaces*

[Setting the Protocol MTU](#)

## mru

### IN THIS SECTION

- [Syntax | 1020](#)
- [Hierarchy Level | 1020](#)
- [Description | 1021](#)
- [Options | 1021](#)
- [Required Privilege Level | 1021](#)
- [Release Information | 1021](#)

## Syntax

```
mru mru;
```

## Hierarchy Level

```
[edit dynamic-profiles name interfaces name gigether-options],
[edit dynamic-profiles name logical-systems name interfaces name gigether-
options],
[edit interfaces name gigether-options]
```

## Description

Configure the maximum receive unit (MRU) of the interface in bytes. The maximum receive unit of an interface indicates the largest size of a packet that the interface can accept. You can configure the parameters so that the value of MRU equals the value of MTU. You can also configure different values for MRU and MTU. When a device receives packets whose size is greater than the interface MRU, those packets are dropped by the device's forwarding plane.

## Options

*mr* MRU size in bytes.

- **Range:** 256 through 16008 bytes.

## Required Privilege Level

interface

## Release Information

Statement introduced in Junos OS Release 19.1R1.

## RELATED DOCUMENTATION

[Ethernet Interfaces Overview | 2](#)

[gigether-options | 896](#)

## multicast-statistics

### IN THIS SECTION

- [Syntax | 1022](#)
- [Hierarchy Level | 1022](#)
- [Description | 1022](#)
- [Default | 1022](#)

- Required Privilege Level | 1022
- Release Information | 1022

## Syntax

```
multicast-statistics;
```

## Hierarchy Level

```
[edit interfaces interface-name]
```

## Description

For Ethernet, SONET, aggregated Ethernet, and aggregated SONET interfaces in T Series or TX Matrix routers, specify support for multicast statistics on a physical interface to enable multicast accounting for all the logical interfaces below the physical interface.

## Default

not enabled—must be configured to enable

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 10.2.

## RELATED DOCUMENTATION

[Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces](#)

[Configuring Multicast Statistics Collection on Aggregated SONET Interfaces](#)

[Configuring Multicast Statistics Collection on Ethernet Interfaces | 9](#)

[Configuring Multicast Statistics Collection on SONET Interfaces](#)

## multiservice

### IN THIS SECTION

- [Syntax | 1023](#)
- [Hierarchy Level | 1024](#)
- [Description | 1024](#)
- [Options | 1024](#)
- [Required Privilege Level | 1024](#)
- [Release Information | 1024](#)

### Syntax

```
multiservice {
 source-mac;
 destination-mac;
 payload {
 ip {
 layer-3 (source-ip-only | destination-ip-only);
 layer-4;
 }
 }
 symmetric-hash {
 complement;
 }
}
```

## Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key family]
```

## Description

(MX Series 5G Universal Routing Platforms only) Configure data used in a hash key for the **multiservice** protocol family when configuring PIC-level symmetrical hashing for load balancing on an 802.3ad Link Aggregation Group.

## Options

- destination-mac** Include destination MAC address in the hash key.
- payload** Include payload data in the hash key. This option has the following suboptions:
- **layer-3**—Include Layer 3 IP information in the hash key.
  - **layer-4**—Include Layer 4 IP information in the hash key.
- source-mac** Include source MAC address in the hash key.
- symmetric-hash** Create a symmetric hash or symmetric hash complement key with any attribute.

## 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 Release 9.6.

## RELATED DOCUMENTATION

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

## negotiate-address

### IN THIS SECTION

- [Syntax | 1025](#)
- [Hierarchy Level | 1025](#)
- [Description | 1025](#)
- [Required Privilege Level | 1025](#)
- [Release Information | 1025](#)

### Syntax

```
negotiate-address;
```

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family inet]
```

### Description

For interfaces with PPP encapsulation, enable the interface to be assigned an IP address by the remote end.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring IPCP Options for Interfaces with PPP Encapsulation](#)

[address](#) | [791](#)

[unnumbered-address \(PPP\)](#) | [1219](#)

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

## negotiation-options

### IN THIS SECTION

- [Syntax](#) | [1026](#)
- [Hierarchy Level](#) | [1026](#)
- [Description](#) | [1026](#)
- [Required Privilege Level](#) | [1027](#)
- [Release Information](#) | [1027](#)

### Syntax

```
negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
}
```

### Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name]
```

### Description

Enable and disable IEEE 802.3ah Operation, Administration, and Management (OAM) features for Ethernet interfaces.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.4.

## RELATED DOCUMENTATION

| *IEEE 802.3ah OAM Link-Fault Management Overview*

## no-adaptive

### IN THIS SECTION

- [Syntax | 1027](#)
- [Hierarchy Level | 1028](#)
- [Description | 1028](#)
- [Required Privilege Level | 1028](#)
- [Release Information | 1028](#)

## Syntax

```
no-adaptive;
```

## Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options load-
balance],
[edit dynamic-profiles name interfaces name logical-tunnel-options load-balance],
[edit dynamic-profiles name interfaces interface-range name aggregated-ether-
options load-balance],
[edit dynamic-profiles name interfaces interface-range name logical-tunnel-
options load-balance],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-
ether-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces name logical-tunnel-
options load-balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name
aggregated-ether-options load-balance],
[edit dynamic-profiles name logical-systems name interfaces interface-range name
logical-tunnel-options load-balance],
[edit interfaces name aggregated-ether-options load-balance],
[edit interfaces name logical-tunnel-options load-balance],
[edit interfaces interface-range name aggregated-ether-options load-
balance],
[edit interfaces interface-range name logical-tunnel-options load-
balance]
```

## Description

Disables the adaptive load-balancing solution configured on the aggregated Ethernet bundle to distribute traffic by using a feedback mechanism.

## Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 13.2R3.

## RELATED DOCUMENTATION

*Understanding Aggregated Ethernet Load Balancing*

## no-allow-link-events

### IN THIS SECTION

- [Syntax | 1029](#)
- [Hierarchy Level | 1029](#)
- [Description | 1029](#)
- [Required Privilege Level | 1029](#)
- [Release Information | 1030](#)

### Syntax

```
no-allow-link-events;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name
negotiation-options]
```

### Description

Disable the sending of link event TLVs.

### 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 Release 8.4.

### RELATED DOCUMENTATION

| *Disabling the Sending of Link Event TLVs*

## no-auto-mdix

### IN THIS SECTION

- [Syntax | 1030](#)
- [Hierarchy Level | 1030](#)
- [Description | 1030](#)
- [Default | 1031](#)
- [Options | 1031](#)
- [Required Privilege Level | 1031](#)
- [Release Information | 1031](#)

### Syntax

```
no-auto-mdix;
```

### Hierarchy Level

```
[edit interface ge-fpc/port/pic together-options]
```

### Description

Disable the Auto MDI/MDIX feature.

MX Series routers with Gigabit Ethernet interfaces automatically detect MDI and MDIX port connections. Use this statement to override the default setting. Remove this statement to return to the default setting.

## Default

Auto MDI/MDIX is enabled by default.

## Options

There are no options for this statement.

## 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 Release 9.5.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

## RELATED DOCUMENTATION

[Ethernet Interfaces Overview | 2](#)

[gigether-options | 896](#)

## no-keepalives

### IN THIS SECTION

- [Syntax | 1032](#)
- [Hierarchy Level | 1032](#)
- [Description | 1032](#)

- Required Privilege Level | 1032
- Release Information | 1032

## Syntax

```
no-keepalives;
```

## Hierarchy Level

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

## Description

Disable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation. The default keepalive interval is 10 seconds.

For ATM2 IQ interfaces only, you can disable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC encapsulation.
- **atm-ppp-vc-mux**—PPP over AAL5 multiplex encapsulation.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Keepalives](#)

[Disabling the Sending of PPPoE Keepalive Messages | 55](#)

[Configuring Frame Relay Keepalives](#)

## no-pre-classifier

### IN THIS SECTION

- [Syntax | 1033](#)
- [Hierarchy Level | 1033](#)
- [Description | 1033](#)
- [Default | 1034](#)
- [Required Privilege Level | 1034](#)
- [Release Information | 1034](#)

### Syntax

```
no-pre-classifier;
```

### Hierarchy Level

```
[edit chassis fpc n pic n]
```

### Description

Specify disabling the control queue for all ports on the 10-Gigabit Ethernet LAN/WAN PIC. Deleting this configuration re-enables the control queue feature on all ports of the 10-Gigabit Ethernet LAN/WAN PIC.

**NOTE:** For the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PD-5-10XGE-SFPP), the control queue has a rate limiter to limit the control traffic to 2 Mbps (fixed, not user-configurable) per port. If the transit control traffic crosses this limit, then it can cause drops on locally terminating control traffic, causing flap of protocols such as BGP and OSPF. To avoid the control traffic being dropped, configure the **no-pre-classifier** statement to disable the control queue.

## Default

The **no-pre-classifier** statement is not configured and the control queue is operational.

## 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 Release 10.4.

## RELATED DOCUMENTATION

[10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 174](#)

[Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 209](#)

## no-send-pads-ac-info

### IN THIS SECTION

- [Syntax | 1035](#)
- [Hierarchy Level | 1035](#)
- [Description | 1035](#)

- Required Privilege Level | 1035
- Release Information | 1035

## Syntax

```
no-send-pads-ac-info;
```

## Hierarchy Level

```
[edit protocols pppoe]
```

## Description

Prevent the router from sending the AC-Name and AC-Cookie tags in the PPPoE Active Discovery Session (PADS) packet. When you configure this statement, it affects PADS packets sent on all PPPoE interfaces configured on the router after the command is issued; it has no effect on previously created PPPoE interfaces. By default, the AC-Name and AC-Cookie tags are transmitted in the PADS packet, along with the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags.

**NOTE:** In Junos OS Release 12.1 and earlier, only the Service-Name, Host-Uniq, Relay-Session-Id, and PPP-Max-Payload tags are contained in the PADS packet by default. The AC-Name and AC-Cookie tags are not transmitted in the PADS packet by default.

## 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 Release 12.2.

## RELATED DOCUMENTATION

| *Disabling the Sending of PPPoE Access Concentrator Tags in PADS Packets*

## no-send-pads-error

### IN THIS SECTION

- [Syntax | 1036](#)
- [Hierarchy Level | 1036](#)
- [Description | 1036](#)
- [Required Privilege Level | 1036](#)
- [Release Information | 1037](#)

### Syntax

```
no-send-pads-error;
```

### Hierarchy Level

```
[edit protocols pppoe]
```

### Description

Discard PADR messages to prevent transmission of PADS control packets with AC-System-Error tags.

### 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 Release 12.3.

### RELATED DOCUMENTATION

| *Discarding PADR Messages to Accommodate Abnormal CPE Behavior*

## non-revertive (Chassis)

### IN THIS SECTION

- [Syntax | 1037](#)
- [Hierarchy Level | 1037](#)
- [Description | 1037](#)
- [Required Privilege Level | 1038](#)
- [Release Information | 1038](#)

### Syntax

```
non-revertive;
```

### Hierarchy Level

```
[edit chassis aggregated-devices ethernet lacp link-protection]
```

### Description

Disable the ability to switch to a better priority link (if one is available) once a link is established as active and a collection or distribution is enabled.

**BEST PRACTICE:** (MX Series) By default, Link Aggregation Control Protocol link protection is revertive. This means that after the current link becomes active, the router switches to a higher-priority link if one becomes operational or is added to the aggregated Ethernet bundle. In a highly scaled configuration over aggregated Ethernet, we recommend that you prevent the router from performing such a switch by including the **non-revertive** statement. Failure to do so may result in some traffic loss if a MIC on which a member interface is located reboots. Using the **non-revertive** statement for this purpose is not effective if both the primary and secondary interfaces are on the MIC that reboots.

## 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 Release 9.3.

## RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## non-revertive (Interfaces)

### IN THIS SECTION

- [Syntax | 1039](#)
- [Hierarchy Level | 1039](#)
- [Description | 1039](#)
- [Required Privilege Level | 1039](#)
- [Release Information | 1039](#)

## Syntax

```
non-revertive;
```

## Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

## Description

Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.

## 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 Release 9.3.

## RELATED DOCUMENTATION

[link-protection](#) | 964

[Configuring Aggregated Ethernet Link Protection](#) | 106

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## number-of-ports

### IN THIS SECTION

- [Syntax | 1040](#)
- [Hierarchy Level | 1040](#)
- [Hierarchy Level | 1040](#)
- [Description | 1040](#)
- [Options | 1041](#)
- [Required Privilege Level | 1041](#)
- [Release Information | 1041](#)

### Syntax

```
number-of-ports number-of-active-physical-ports;
```

### Hierarchy Level

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

### Hierarchy Level

```
[edit chassis fpc fpc-slot pic pic-number pic-mode pic-speed]
```

### Description

Administratively enable physical ports, for example, to prevent oversubscription of the line card fabric interface. By default, all available ports are enabled. When disabled, the LED on the affected line card will appear yellow on capable line cards.

(MX Series with 16x10GE MPC, MPC3, MPC4, MPC5, and MPC6) You can disable a subset of the physical ports available on the Packet Forwarding Engines of the 16x10GE MPC, and for MICs installed in MPC3, MPC4, MPC5, and MPC6. Specify either 8 or 12 ports by using this statement. When eight active ports are configured, two ports per Packet Forwarding Engine are disabled, and the LEDs on the MPC appear **yellow**. When you specify 12 active ports, one port per Packet Forwarding Engine is disabled and the corresponding LED appear **yellow**. When you do not include this statement in the configuration, all 16 default ports on the MPC are active.

(MX Series with MPC7E-MRATE, MPC8E, and MPC9E) To ensure guaranteed bandwidth by preventing fabric oversubscription, you can disable a subset of the physical ports available on MPC7E-MRATE, MPC8E, and MPC9E. For information about the active ports for MPC7E-MRATE, MPC8E, and MPC9E, see [Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription](#).

(MX204 routers) To ensure guaranteed bandwidth by preventing oversubscription, you can disable a subset of the physical ports available on MX204 routers. For information about the active ports for MX204 routers, see [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router](#)

(EX9200 switches)

## Options

*number-of-active-physical-ports*—Specify the number of physical ports to enable on PICs or MICs on an MPC.

## 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 Release 10.1.

Support for MPC3, MPC4, MPC5, and MPC6 introduced in Junos OS Release 13.3R2.

Support for MPC7E-MRATE MPC introduced in Junos OS Release 15.1F4.

Support for MPC8E and MPC9E introduced in Junos OS Release 15.1F5.

Support for MX10003 MPC introduced in Junos OS Release 17.3R1

Support for MX204 routers introduced in Junos OS Release 17.4R1

Statement introduced in Junos OS Release 16.1 for EX9200 switches.

## RELATED DOCUMENTATION

[Configuring the Number of Active Ports on 16x10GE 3D MPC](#)

[Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription](#)

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

[Introduction to Port Speed | 370](#)

[Introduction to Port Speed | 370](#)

## number-of-sub-ports

### IN THIS SECTION

- [Syntax | 1042](#)
- [Hierarchy Level | 1042](#)
- [Description | 1042](#)
- [Options | 1044](#)
- [Required Privilege Level | 1044](#)
- [Release Information | 1044](#)

### Syntax

```
number-of-sub-ports <number-of-sub-ports>;
```

### Hierarchy Level

```
[edit chassis fpc fpc-slot pic pic-number port port-num]
[edit interfaces interface-name]
```

### Description

*For PTX10003-80C, PTX10003-160C router:*

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.

*For MPC10E-15C-MRATE supported on MX240, MX480, MX960 routers:*

To configure the number of sub-channels for a particular port if the optics are used in a channelized mode.

(Channelized mode) To specify the number of IFDs (or interfaces) that need to be created on a physical port for a specified speed, use the **number-of-sub-ports <number-of-sub-ports>** configuration statement. For example, on a given port that supports 4x10GE mode, if the **number-of-sub-ports** to 2, then two IFDs are created, namely et-x/y/z:0 and et-x/y/z:1.

The default value of channelized 10-Gigabit Ethernet interface is 4. The number of sub-ports that can be configured are, 1, 2, 3, or 4. You must set the **number-of-sub-ports** to 4, to channelize 40-Gigabit Ethernet interface to four 10-Gigabit Ethernet interfaces.

The **number-of-sub-ports** configuration statement can be used with rate selectability configuration at both PIC level and port level. This configuration statement is effective only when the port speed is 10 Gbps.

**NOTE:** You can configure the *number-of-sub-ports* only for 10-Gbps speed. For other speeds, this configuration is not supported.

(MPC11E) To specify the number of interfaces to be created on a physical port.

(JNP10K-LC1201 and JNP10K-LC1202) Specifies the number of channelized interfaces that you can configure on a physical port with the specified speed. The **edit chassis** hierarchy is available for 20.2R1

Release on PTX10001-36MR. From 20.3R1 onwards you must use only the **edit interface***interface-name* hierarchy.

## Options

**number-of-sub-ports** *number-of-subports* Specify the number of sub-ports per physical port. For PTX10003-80C and PTX10003-160C routers, the values are 1, 2, and 4. On MPC10E-15C-MRATE line card, the values are 1, 2, 3, and 4. On JNP10K-LC1201, the values are 0 through 7. On JNP10K-LC1202, the values are 1 through 8.

## 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 Release 19.1R1.

Interface hierarchy introduced in Junos OS Evolved Release 20.1R2 for JNP10K-LC1201 line cards on PTX10008 routers.

## RELATED DOCUMENTATION

[speed \(Ethernet\) | 1137](#)

[Introduction to Port Speed | 370](#)

[Configuring the Port Speed on the JNP10K-LC1201 by Using New Port Profile Configuration | 452](#)

## oam

### IN THIS SECTION

● [Syntax | 1045](#)

● [Hierarchy Level | 1047](#)

- Description | 1048
- Required Privilege Level | 1048
- Release Information | 1048

## Syntax

```
oam {
 ethernet {
 connectivity-fault-management {
 action-profile profile-name {
 default-actions {
 interface-down;
 }
 }
 }
 performance-monitoring {
 delegate-server-processing;
 hardware-assisted-timestamping;
 hardware-assisted-keepalives;
 sla-iterator-profiles {
 profile-name {
 avg-fd-twoway-threshold;
 avg-ifdv-twoway-threshold;
 avg-flr-forward-threshold;
 avg-flr-backward-threshold;
 disable;
 calculation-weight {
 delay delay-weight;
 delay-variation delay-variation-weight;
 }
 cycle-time milliseconds;
 iteration-period connections;
 measurement-type (loss | statistical-frame-loss | two-
way-delay);
 }
 }
 }
 }
 linktrace {
 age (30m | 10m | 1m | 30s | 10s);
 }
}
```

```

 path-database-size path-database-size;
 }
 maintenance-domain domain-name {
 level number;
 name-format (character-string | none | dns | mac+2octet);
 maintenance-association ma-name {
 short-name-format (character-string | vlan | 2octet |
rfc-2685-vpn-id);
 protect-maintenance-association protect-ma-name;
 remote-maintenance-association remote-ma-name;
 continuity-check {
 convey-loss-threshold;
 hold-interval minutes;
 interface-status-tlv;
 interval (100ms | 10m | 10ms | 10s | 1m | 1s);
 loss-threshold number;
 port-status-tlv;
 }
 mep mep-id {
 auto-discovery;
 direction (up | down);
 interface interface-name (protect | working);
 lowest-priority-defect (all-defects | err-xcon | mac-rem-
err-xcon | no-defect | rem-err-xcon | xcon);
 priority number;
 remote-mep mep-id {
 action-profile profile-name;
 sla-iterator-profile profile-name {
 data-tlv-size size;
 iteration-count count-value;
 priority priority-value;
 }
 }
 }
 }
 }
}
link-fault-management {
 action-profile profile-name {
 action {
 link-down;
 send-critical-event;
 syslog;

```



## Description

For Ethernet interfaces on M320, M120, MX Series, and T Series routers and PTX Series Packet Transport Routers, provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.2.

## RELATED DOCUMENTATION

| [IEEE 802.3ah OAM Link-Fault Management Overview](#)

## optics-options

### IN THIS SECTION

- [Syntax | 1049](#)
- [Hierarchy Level | 1049](#)
- [Description | 1050](#)
- [Options | 1050](#)
- [Required Privilege Level | 1050](#)
- [Release Information | 1050](#)

## Syntax

```

optics-options {
 alarm low-light-alarm {
 (link-down | syslog);
 }
 encoding (differential | non-differential);
 fec (hgfec | sdfec | sdfec15 | sdfec25);
 high-polarization;
 host-side-fec (off | on | not-supported);
 (is-ma| no-is-ma);
 (lane| lane-all);
 (laser-enable | no-laser-enable);
 loopback;
 los-alarm-threshold dBm;
 los-warning-threshold dBm;
 low-power-mode;
 media-side-fec (off | on | unsupported);
 modulation-format (16qam | 8qam | qpsk);
 signal-degrade {
 ber-threshold-clear ber-threshold-clear;
 ber-threshold-signal-degrade ber-threshold-signal-degrade;
 interval milliseconds;
 q-threshold-signal-degrade dB;
 q-threshold-signal-degrade-clear dB;
 }
 tca tca-identifier (enable-tca | no-enable-tca) (threshold number |
threshold-24hrs number);
 tx-power dbm;
 warning low-light-warning {
 (link-down | syslog);
 }
 wavelength nm;
 loopback;
}

```

## Hierarchy Level

```
[edit interfaces interface-name]
```

## Description

For 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces only, configure full C-band International Telecommunication Union (ITU)-Grid tunable optics.

On the PTX Series routers, when an interface is configured in 8QAM mode, you must configure both the optics from a AC400 module with the same optics-options for the links to come up.

## Options

- **host-side-fec**—Enable or disable the FEC on the host side optics.
- **loopback**—Displays the electrical loopback status of QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.
- **los-alarm-threshold**—Specify the Loss of Signal (LOS) threshold above which an alarm is raised.
- **los-warning-threshold**—Specify the Loss of Signal (LOS) threshold above which a warning is issued.
- **low-power-mode**—Forces the optics to be on low power mode. By default, it is disabled. On channelized interfaces, you can configure the **low-power-mode** on channel 0 only.
- **media-side-fec**—Enable or disable the FEC on media side optics.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

**alarm** option and **warning** options introduced in Junos OS Release 10.0.

Statement and **tx-power** option introduced in Junos OS Release 13.2 for PTX Series routers.

**tca** option introduced in Junos OS Release 14.2 for PTX Series routers.

**los-alarm-threshold** and **los-alarm-warning** options introduced in Junos OS Release 14.2.

**loopback** option introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

**lane**, **lane-allhost-side-fec**, and **media-side-fec** options introduced in Junos OS Evolved Release 21.1R1

## RELATED DOCUMENTATION

[Ethernet DWDM Interface Wavelength Overview | 552](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

[Supported Forward Error Correction Modes on ACX6360 Router | 550](#)

[lane | 944](#)

[lane-all | 947](#)

## otn-options

### IN THIS SECTION

- [Syntax | 1051](#)
- [Hierarchy Level | 1052](#)
- [Description | 1053](#)
- [Options | 1053](#)
- [Required Privilege Level | 1053](#)
- [Release Information | 1053](#)

## Syntax

```
otn-options {
 bytes (otn-options) transmit-payload-type value;
 fec (efec | gfec | gfec-sdfec | none);
 (is-ma | no-is-ma);
 (laser-enable | no-laser-enable);
 (line-loopback | no-line-loopback);
 (local-loopback | no-local-loopback);
 (odu-ttim-action-enable | no-odu-ttim-action-enable);
 (otu-ttim-action-enable | no-otu-ttim-action-enable);
 odu-delay-management {
```

```

 (bypass | no-bypass);
 (monitor-end-point | no-monitor-end-point);
 number-of-frames value;
 (no-start-measurement | start-measurement);
}
odu-signal-degrade {
 ber-threshold-clear value;
 ber-threshold-signal-degrade value;
 interval value;
}
(prbs | no-prbs);
preemptive-fast-reroute {
 (backward-frr-enable | no-backward-frr-enable);
 (signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
 odu-backward-frr-enable | no-odu-backward-frr-enable;
 odu-signal-degrade-monitor-enable | no-odu-signal-degrade-monitor-enable;
}
rate {
 (fixed-stuff-bytes | no-fixed-stuff-bytes);
 oc192;
 otu4;
 (pass-through | no-pass-through);
}
signal-degrade {
 ber-threshold-clear value;
 ber-threshold-signal-degrade value;
 interval value;
}
tca tca-identifier (enable-tca | no-enable-tca) (threshold number |
threshold-24hrs number);
transport-monitoring;
trigger trigger-identifier;
tti tti-identifier;
}

```

## Hierarchy Level

```

[edit interfaces ge-fpc/pic/port]
[edit interfaces xe-fpc/pic/port]
[edit interfaces et-fpc/pic/port]

```

## Description

Specify the Ethernet optical transport network (OTN) interface and options.

## Options

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

## Required Privilege Level

interfaces—To view this statement in the configuration.

interfaces-control—To add this statement to the configuration.

## Release Information

Statement introduced in Junos OS Release 9.4.

**bytes**, **is-ma**, **local-loopback**, **no-is-ma**, **no-local-loopback**, **no-odu-ttim-action-enable**, **no-otu-ttim-action-enable**, **no-prbs**, **odu-delay-management**, **odu-ttim-action-enable**, **otu-ttim-action-enable**, **prbs**, **preemptive-fast-reroute**, and **signal-degrade** statements introduced in Junos OS Release 13.2 for PTX Series routers.

**oc192** statement introduced in Junos OS Release 13.3R3 for MX Series routers.

**odu-signal-degrade**, **odu-backward-frr-enable** | **no-odu-backward-frr-enable**, **odu-signal-degrade-monitor-enable** | **no-odu-signal-degrade-monitor-enable** statements introduced in Junos OS Release 14.1R2 and 14.2 for P2-100GE-OTN PIC in PTX5000 routers.

**tca** option introduced in Junos OS Release 14.2 for PTX Series routers.

**bytes**, **line-loopback**, **local-loopback**, **preemptive-fast-reroute**, **tca**, **trigger**, **prbs**, and **tti** statements introduced in 18.3R1 for ACX6360 routers.

## RELATED DOCUMENTATION

---

[10-Gigabit Ethernet OTN Options Configuration Overview](#) | 459

---

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 459

---

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 561

## output-policer

### IN THIS SECTION

- [Syntax | 1054](#)
- [Hierarchy Level | 1054](#)
- [Description | 1054](#)
- [Options | 1054](#)
- [Required Privilege Level | 1055](#)
- [Release Information | 1055](#)

### Syntax

```
output-policer policer-name;
```

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number layer2-policer]
```

### Description

Apply a single-rate two-color policer to the Layer 2 output traffic at the logical interface. The **output-policer** and **output-three-color** statements are mutually exclusive.

### Options

***policer-name***—Name of the single-rate two-color policer that you define at the **[edit firewall]** hierarchy level.

## 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 Release 8.2.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 271](#)

*input-policer*

*input-three-color*

*layer2-policer*

*logical-interface-policer*

*output-three-color*

## output-priority-map

### IN THIS SECTION

- [Syntax | 1056](#)
- [Hierarchy Level | 1056](#)
- [Description | 1056](#)
- [Required Privilege Level | 1056](#)
- [Release Information | 1056](#)

## Syntax

```
output-priority-map {
 classifier {
 premium {
 forwarding-class class-name {
 loss-priority (high | low);
 }
 }
 }
}
```

## Hierarchy Level

```
[edit interfaces interface-name together-options ethernet-switch-profile
ethernet-policer-profile]
[edit interfaces interface-name ether-options ethernet-switch-profile ethernet-
policer-profile]
```

## Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the output policer priority map to be applied to outgoing frames on this interface.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[input-priority-map | 920](#)

## output-three-color

### IN THIS SECTION

- [Syntax | 1057](#)
- [Hierarchy Level | 1057](#)
- [Description | 1057](#)
- [Options | 1058](#)
- [Required Privilege Level | 1058](#)
- [Release Information | 1058](#)

### Syntax

```
output-three-color policer-name;
```

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number layer2-policer]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number layer2-policer]
```

### Description

Apply a single-rate or two-rate three-color policer to the Layer 2 output traffic at the logical interface. The **output-three-color** and **output-policer** statements are mutually exclusive.

## Options

*policer-name*—Name of the single-rate or two-rate three-color policer.

## 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 Release 8.2.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 271](#)

*input-three-color*

*input-policer*

*layer2-policer*

*logical-interface-policer*

*output-policer*

## output-vlan-map (Aggregated Ethernet)

### IN THIS SECTION

- [Syntax | 1059](#)
- [Hierarchy Level | 1059](#)
- [Description | 1059](#)
- [Required Privilege Level | 1059](#)
- [Release Information | 1059](#)

## Syntax

```
output-vlan-map {
 (pop | push | swap);
 tag-protocol-id tpid;
 vlan-id number;
}
```

## Hierarchy Level

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

## Description

Define the rewrite profile to be applied to outgoing frames on this logical interface. On MX Series routers, this statement only applies to aggregated Ethernet interfaces using Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and 100-Gigabit Ethernet Type 5 PIC with CFP.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.2.

Starting in Junos OS Release 17.3R1, input-vlan-map for outer vlan is supported for L2 circuit over aggregated Ethernet interfaces for QFX10000 Series switches.

## RELATED DOCUMENTATION

[Stacking and Rewriting Gigabit Ethernet VLAN Tags](#)

[input-vlan-map \(Aggregated Ethernet\) | 923](#)

## pado-advertise

### IN THIS SECTION

- [Syntax | 1060](#)
- [Hierarchy Level | 1060](#)
- [Description | 1060](#)
- [Required Privilege Level | 1060](#)
- [Release Information | 1061](#)

### Syntax

```
pado-advertise;
```

### Hierarchy Level

```
[edit protocols pppoe]
```

### Description

Enable named services configured in PPPoE service name tables to be advertised in PPPoE Active Discovery Offer (PADO) control packets. By default, advertisement of named services in PADO packets is disabled.

**NOTE:** If you enable advertisement of named services in PADO packets, make sure the number and length of all advertised service entries does not exceed the maximum transmission unit (MTU) size of the PPPoE underlying interface.

### 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 Release 10.2.

### RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

*Enabling Advertisement of Named Services in PADO Control Packets*

## passive-monitor-mode

### IN THIS SECTION

- [Syntax | 1061](#)
- [Hierarchy Level | 1061](#)
- [Description | 1062](#)
- [Required Privilege Level | 1062](#)
- [Release Information | 1062](#)

### Syntax

```
passive-monitor-mode;
```

### Hierarchy Level

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

## Description

Monitor packet flows from another router. If you include this statement in the configuration, the interface does not send keepalives or alarms, and does not participate actively on the network.

This statement is supported on ATM, Ethernet, and SONET/SDH interfaces. For more information, see [ATM Interfaces User Guide for Routing Devices](#).

For ATM and Ethernet interfaces, you can include this statement on the physical interface only.

For SONET/SDH interfaces, you can include this statement on the logical interface only.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Enabling Passive Monitoring on ATM Interfaces](#)

[Passive Monitoring on Ethernet Interfaces Overview | 595](#)

[Enabling Packet Flow Monitoring on SONET/SDH Interfaces](#)

*multiservice-options*

[Junos OS Services Interfaces Library for Routing Devices](#)

## payload

### IN THIS SECTION

- [Syntax | 1063](#)
- [Hierarchy Level | 1063](#)
- [Description | 1063](#)

- Options | 1063
- Required Privilege Level | 1063
- Release Information | 1064

## Syntax

```
payload {
 ip {
 layer-3;
 layer-4;
 }
}
```

## Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key family multiservice]
```

## Description

(MX Series 5G Universal Routing Platforms only) Include payload data in a hash key for the **multiservice** protocol family when configuring PIC-level symmetrical load balancing on an 802.3ad Link Aggregation Group.

## Options

**ip**—Include IPv4 payload data in the hash key. This option has the following suboptions:

- **layer-3**—Include Layer 3 IP information in the hash key.
- **layer-4**—Include Layer 4 IP information in the hash key.

## 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 Release 9.6.

### RELATED DOCUMENTATION

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

## pdu-interval

### IN THIS SECTION

- [Syntax | 1064](#)
- [Hierarchy Level | 1064](#)
- [Description | 1065](#)
- [Options | 1065](#)
- [Required Privilege Level | 1065](#)
- [Release Information | 1065](#)

### Syntax

```
pdu-interval interval;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

## Description

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

## Options

*interval*—Periodic OAM PDU sending interval.

- **Range:** For MX, M, T, ACX, Series routers, SRX Series firewalls and EX Series switches – 100 through 1000 milliseconds
- **Default:** 1000 milliseconds

## Required Privilege Level

interface—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.2.

## RELATED DOCUMENTATION

*Configuring the OAM PDU Interval*

*Example: Configuring Ethernet OAM Link Fault Management*

*Configuring Ethernet OAM Link Fault Management*

## pdu-threshold

### IN THIS SECTION

- [Syntax | 1066](#)
- [Hierarchy Level | 1066](#)
- [Description | 1066](#)
- [Options | 1066](#)
- [Required Privilege Level | 1067](#)
- [Release Information | 1067](#)

### Syntax

```
pdu-threshold threshold-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

### Description

Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

### Options

***threshold-value***—The number of PDUs missed before declaring the peer lost.

- **Range:** 3 through 10 PDUs
- **Default:** 3 PDUs

## 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 Release 8.2.

## RELATED DOCUMENTATION

*Configuring the OAM PDU Threshold*

*Configuring Ethernet OAM Link Fault Management*

## per-flow (Aggregated Ethernet Interfaces)

### IN THIS SECTION

- [Syntax | 1067](#)
- [Hierarchy Level | 1068](#)
- [Description | 1068](#)
- [Options | 1068](#)
- [Required Privilege Level | 1068](#)
- [Release Information | 1068](#)

## Syntax

```
per-flow;
```

## Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful]
```

## Description

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. When multiple flows are transmitted out of an **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.

## Options

**per-flow** Enable the stateful load-distribution mechanism per traffic flow on an aggregated Ethernet interface.

## 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 Release 13.2R1.

## RELATED DOCUMENTATION

| *Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces*

## periodic

### IN THIS SECTION

- [Syntax \(EX Series\) | 1069](#)
- [Syntax \(QFX Series\) | 1069](#)
- [Hierarchy Level \(EX Series\) | 1069](#)
- [Hierarchy Level \(QFX Series\) | 1069](#)
- [Description | 1070](#)
- [Options | 1070](#)
- [Required Privilege Level | 1070](#)
- [Release Information | 1070](#)

### Syntax (EX Series)

```
periodic interval;
```

### Syntax (QFX Series)

```
periodic (fast | slow);
```

### Hierarchy Level (EX Series)

```
[edit interfaces aex aggregated-ether-options lacp],
[edit interfaces interface-range name aggregated-ether-options lacp]
```

### Hierarchy Level (QFX Series)

```
[edit interfaces aex aggregated-ether-options lacp]
```

## Description

For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

## Options

*interval*—Interval for periodic transmission of LACP packets.

- **fast**—Transmit packets every second.
- **slow**—Transmit packets every 30 seconds.
- **Default: fast**

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#)

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

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

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

*Understanding Aggregated Ethernet Interfaces and LACP for Switches*

[Junos OS Network Interfaces Library for Routing Devices](#)

## pic-mode

### IN THIS SECTION

- [Syntax | 1071](#)
- [Hierarchy Level | 1071](#)
- [Description | 1071](#)
- [Options | 1072](#)
- [Required Privilege Level | 1072](#)
- [Release Information | 1073](#)

### Syntax

```
pic-mode pic-speed;
```

### Hierarchy Level

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

### Description

Configure the operating speed of all ports on the MPC7E-MRATE MPC, MIC-MRATE MIC, MX10003 MPC, and MX204 routers.

(MX240, MX480, MX960, MX2010, and MX2020 routers with MPC7E-MRATE) To configure 100 Gbps, 10 Gbps, and 40 Gbps speed on all supported ports, specify **100G**, **10G**, or **40G**, respectively, as the speed for the specified PIC. All the six ports of PIC 0 and PIC 1 of an MPC7E-MRATE MPC support 10-Gbps and 40-Gbps speeds. However, only ports 2 and 5 of PIC 0 and PIC 1 support 100-Gbps speed. Therefore, if you configure **100G** as the operating speed of a PIC, the PIC is rebooted and the ports 0, 1, 3, and 4 are disabled.

(MX2010 and MX2020 routers with MIC-MRATE on MPC8E and MPC9E) To configure the port speed at MIC level or on all supported ports, specify **100G**, **10G**, or **40G**, respectively, as the speed for the MIC-MRATE MIC on MPC8E and MPC9E. All the twelve ports of MIC-MRATE MIC support 10 Gbps and 40 Gbps speeds. When you configure the port speed as 100 Gbps at the PIC level for MPC8E, you

can configure only 4 ports of the 12 MIC-MRATE ports on MPC8E to operate at 100 Gbps port speed. The other ports are disabled. Therefore, if you configure **100G** as the operating speed for ports 0, 1, 6, and 7, the other ports are disabled on MPC8E. Similarly, when you configure the port speed as 100 Gbps at the PIC level you can configure only 8 ports of the 12 MIC-MRATE ports on MPC9E with 100 Gbps port speed. Therefore, if you configure **100G** as the operating speed for ports 0, 1, 2, 3, 6, 7, 8, and 9, the other ports can support only 40 Gbps or 10 Gbps. However, enabling port speed of 40 Gbps or 10 Gbps at the PIC level, enables all ports and sets the desired port speed on all ports.

(MX10003 routers with MX10003 MPC) To configure 100 Gbps, 10 Gbps, and 40 Gbps speed on all supported ports, specify 100G, 10G, or 40G, respectively, as the speed for the specified PIC. All the six ports of the fixed port PIC support 10-Gbps and 40-Gbps speeds. All the 12 ports of the Multi-rate MIC support 100-Gbps, 10-Gbps and 40-Gbps speeds. To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the **pic-mode** statement and specify the port speed. For more information see [MX10003 MPC \(Multi-Rate\)](#) and [Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription on MX10003 MPC](#).

(MX204 routers) To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the **pic-mode** statement and specify the port speed. The MX204 has four rate-selectable ports (referred to a PIC 0 ports) that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet port, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The MX204 also has eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports).

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 [page 1040](#) on PIC 1 must be configured to 0 only. For more information, see [MX204 Router Overview](#) and [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router](#).

## Options

**pic-speed**—Operating speed of the interfaces configured on the ports of an MPC7E-MRATE MPC.

**100G**—Supported ports operate at 100 Gbps speed.

**10G**—Supported ports operate at 10 Gbps speed.

**40G**—Supported ports operate at 40 Gbps speed.

- **Default: 10G**

## 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 Release 15.1F4.

Statement introduced in Junos OS Release 17.3R1 for MX10003 routers with the MX10003 MPC.

## RELATED DOCUMENTATION

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

[Introduction to Port Speed | 370](#)

[Introduction to Port Speed | 370](#)

## policer (CFM Firewall)

### IN THIS SECTION

- [Syntax | 1073](#)
- [Hierarchy Level | 1074](#)
- [Description | 1074](#)
- [Required Privilege Level | 1074](#)
- [Release Information | 1074](#)

## Syntax

```
policer cfm-policer {
 if-exceeding {
 bandwidth-limit 8k;
 burst-size-limit 2k;
 }
}
```

```
 then discard;
}
```

## Hierarchy Level

```
[edit firewall]
```

## Description

Attach an explicit policer to CFM sessions.

## 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 Release 10.0.

## RELATED DOCUMENTATION

*Configuring Rate Limiting of Ethernet OAM Messages*

*policer (CFM Global)*

*policer (CFM Session)*

## policer (CoS)

### IN THIS SECTION

- [Syntax | 1075](#)
- [Hierarchy Level | 1075](#)
- [Description | 1075](#)

- Options | 1075
- Required Privilege Level | 1076
- Release Information | 1076

## Syntax

```
policer cos-policer-name {
 aggregate {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
 premium {
 bandwidth-limit bps;
 burst-size-limit bytes;
 }
}
```

## Hierarchy Level

```
[edit interfaces interface-name together-options ethernet-switch-profile
ethernet-policer-profile]
```

## Description

For Gigabit Ethernet IQ , Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and 100-Gigabit Ethernet Type 5 PIC with CFP, define a CoS policer template to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits. The premium policer is not supported on MX Series routers or for Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

## Options

***cos-policer-name***—Name of one policer to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Configuring Gigabit Ethernet Policers](#) | 271

## policer (MAC)

### IN THIS SECTION

- [Syntax](#) | 1076
- [Hierarchy Level](#) | 1077
- [Description](#) | 1077
- [Options](#) | 1077
- [Required Privilege Level](#) | 1077
- [Release Information](#) | 1078

## Syntax

```
policer {
 input cos-policer-name;
 output cos-policer-name;
}
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number accept-source-mac mac-address mac-address],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number accept-source-mac mac-address mac-address]
```

## Description

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 100-Gigabit Ethernet Type 5 PIC with CFP, configure MAC policing.

**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.

## Options

**input *cos-policer-name***—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

**output *cos-policer-name***—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

| [Configuring Gigabit Ethernet Policers](#) | 271

## port-priority

### IN THIS SECTION

- [Syntax](#) | 1078
- [Hierarchy Level](#) | 1078
- [Description](#) | 1078
- [Options](#) | 1079
- [Required Privilege Level](#) | 1079
- [Release Information](#) | 1079

### Syntax

```
port-priority priority;
```

### Hierarchy Level

```
[edit interfaces interface-name together-options 802.3ad lacp]
```

### Description

Define LACP port priority at the interface level.

## Options

*priority*—Priority for being elected to be the active port and both collect and distribute traffic. A smaller value indicates a higher priority for being elected.

- **Range:** 0 through 65535
- **Default:** 127

## 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 Release 9.3.

## RELATED DOCUMENTATION

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

## pp0 (Dynamic PPPoE)

### IN THIS SECTION

- [Syntax | 1080](#)
- [Hierarchy Level | 1081](#)
- [Description | 1081](#)
- [Required Privilege Level | 1081](#)
- [Release Information | 1081](#)

## Syntax

```

pp0 {
 unit logical-unit-number {
 keepalives interval seconds;
 no-keepalives;
 pppoe-options {
 underlying-interface interface-name;
 server;
 }
 ppp-options {
 aaa-options aaa-options-name;
 authentication [authentication-protocols];
 chap {
 challenge-length minimum minimum-length maximum maximum-length;
 }
 ignore-magic-number-mismatch;
 initiate-ncp (ip | ipv6 | dual-stack-passive)
 ipcp-suggest-dns-option;
 mru size;
 mtu (size | use-lower-layer);
 on-demand-ip-address;
 pap;
 peer-ip-address-optional;
 }
 family inet {
 unnumbered-address interface-name;
 address address;
 service {
 input {
 service-set service-set-name {
 service-filter filter-name;
 }
 post-service-filter filter-name;
 }
 output {
 service-set service-set-name {
 service-filter filter-name;
 }
 }
 }
 filter {

```



## ppm (Ethernet Switching)

### IN THIS SECTION

- [Syntax | 1082](#)
- [Hierarchy Level | 1082](#)
- [Description | 1082](#)
- [Default | 1083](#)
- [Required Privilege Level | 1083](#)
- [Release Information | 1083](#)

### Syntax

```
ppm {
 centralized;
}
```

### Hierarchy Level

```
[edit protocols lacp]
```

### Description

Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.

This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** configuration statement in the [\[edit routing-options ppm\]](#) statement hierarchy.

## Default

Distributed PPM processing is enabled for all packets that use PPM.

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.4.

## RELATED DOCUMENTATION

[Configuring Distributed Periodic Packet Management on an EX Series Switch \(CLI Procedure\)](#)

[Configuring Distributed Periodic Packet Management](#)

## pppoe-options

### IN THIS SECTION

- [Syntax | 1083](#)
- [Hierarchy Level | 1084](#)
- [Description | 1084](#)
- [Required Privilege Level | 1085](#)
- [Release Information | 1085](#)

## Syntax

```
pppoe-options {
 access-concentrator name;
 auto-reconnect seconds;
```

```

(client | server);
service-name name;
underlying-interface interface-name;
ppp-max-payload ppp-max-payload
}

```

## Hierarchy Level

```

[edit interfaces pp0 unit logical-unit-number],
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-
number]
[set interface ppp interfaceunit logical-unit-number ppp-max-payload ppp-max-
payload],

```

## Description

Configure PPP over Ethernet-specific interface properties.

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

The maximum payload allowed on an Ethernet frame is 1500 bytes. For a PPPoE interface, the PPPoE header uses 6 bytes and the PPP protocol ID uses 2 bytes. This restricts the maximum MTU size on a PPPoE interface to 1492 bytes, which can cause frequent fragmentation and reassembly of larger PPP packets received over the PPPoE interface. To prevent frequent fragmentation and reassembly for PPP packets over Ethernet, you can configure the maximum transmission unit (MTU) and MRU sizes for PPP subscribers.

For PPPoE subscribers, the PPP MRU or PPP MTU size can be greater than 1492 bytes if the PPP-Max-Payload tag is received in the PPPoE Active Discovery Request (PADR) packets.

The PPP-Max-Payload option allows you to override the default behavior of the PPPoE client by providing a maximum size that the PPP payload can support in both sending and receiving directions. The PPPoE server might allow the negotiation of an MRU larger than 1492 octets and the ability to use an MTU larger than 1500 octets.

It is important to set an appropriate value for the MTU size of the physical interface before setting **ppp-max-payload**. The value of **mtu** must be greater than the value of **ppp-max-payload**.

To enable Jumbo frames refer [Understanding Jumbo Frames Support for Ethernet Interfaces](#).

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

**client** Statement introduced in Junos OS Release 8.5.

**server** Statement introduced in Junos OS Release 8.5.

**ppp-max-payload** Statement introduced in Junos OS Release 15.1X49-D100.

## RELATED DOCUMENTATION

| [Configuring a PPPoE Interface](#)

## pppoe-underlying-options (Static and Dynamic Subscribers)

### IN THIS SECTION

- [Syntax | 1085](#)
- [Hierarchy Level | 1086](#)
- [Description | 1086](#)
- [Required Privilege Level | 1086](#)
- [Release Information | 1086](#)

## Syntax

```
pppoe-underlying-options {
 access-concentrator name;
 dynamic-profile profile-name;
```

```

direct-connect
duplicate-protection;
max-sessions number;
max-sessions-vs-a-ignore;
service-name-table table-name;
short-cycle-protection <lockout-time-min minimum-seconds> <lockout-time-max
maximum-seconds> <filter [aci]>;
}

```

## Hierarchy Level

```

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

```

## Description

Configure PPPoE-specific interface properties for the underlying interface on which the router creates a static or dynamic PPPoE logical interface. The underlying interface must be configured with PPPoE (**ppp-over-ether**) encapsulation.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 10.0.

## RELATED DOCUMENTATION

[Configuring PPPoE](#)

*Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*

*Assigning a Service Name Table to a PPPoE Underlying Interface*

## preferred-source-address

### IN THIS SECTION

- [Syntax | 1087](#)
- [Hierarchy Level | 1087](#)
- [Description | 1087](#)
- [Options | 1088](#)
- [Required Privilege Level | 1088](#)
- [Release Information | 1088](#)

### Syntax

```
preferred-source-address address;
```

### Hierarchy Level

```
[edit dynamic-profiles interfaces interface-name unit logical-unit-number family family unnumbered-address interface-name],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family],
```

### Description

For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. To configure the preferred source address dynamically, instead of using this statement, you must include the **\$junos-preferred-source-address** predefined variable for IPv4 (family inet) addresses or the **\$junos-preferred-source-ipv6-address** predefined variable for IPv6 (family inet6) addresses.

Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for IPv4 and IPv6 address families.

**NOTE:** When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the **\$junos-routing-instance** predefined variable, you must not configure a preferred source address, whether with the **\$junos-preferred-source-address** predefined variable, the **\$junos-preferred-source-ipv6-address** predefined variable, or the **preferred-source-address** statement. Configuring the preferred source address in this circumstance causes a commit failure.

## Options

**address**—Secondary IP address of the donor loopback interface. Alternatively, use the **\$junos-preferred-source-address** or the **\$junos-preferred-source-ipv6-address** predefined variable to dynamically apply a preferred source address to the unnumbered Ethernet interface.

## 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 Release 9.2.

Support for the **\$junos-preferred-source-address** and **\$junos-preferred-source-ipv6-address** predefined variables introduced in Junos OS Release 9.6.

## RELATED DOCUMENTATION

---

[Configuring an Unnumbered Interface](#)

---

[Junos OS Network Interfaces Library for Routing Devices](#)

---

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

## premium (Output Priority Map)

### IN THIS SECTION

- [Syntax | 1089](#)
- [Hierarchy Level | 1089](#)
- [Description | 1089](#)
- [Required Privilege Level | 1089](#)
- [Release Information | 1090](#)

### Syntax

```
premium {
 forwarding-class class-name {
 loss-priority (high | low);
 }
}
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile output-priority-map classifier]
```

### Description

For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[input-priority-map | 920](#)

## premium (Policer)

### IN THIS SECTION

- [Syntax | 1090](#)
- [Hierarchy Level | 1090](#)
- [Description | 1091](#)
- [Required Privilege Level | 1091](#)
- [Release Information | 1091](#)

### Syntax

```
premium {
 bandwidth-limit bps;
 burst-size-limit bytes;
}
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile policer cos-policer-name]
```

## Description

Define a policer to apply to nonpremium traffic.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 271

[aggregate \(Gigabit Ethernet CoS Policer\)](#) | 798

[ieee802.1p](#) | 908

## protocol-down

### IN THIS SECTION

- [Syntax](#) | 1092
- [Hierarchy Level](#) | 1092
- [Description](#) | 1092
- [Required Privilege Level](#) | 1092
- [Release Information](#) | 1092

## Syntax

```
protocol-down;
```

## Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event]
```

## Description

Upper layer indication of protocol down event. When the **protocol-down** statement is included, the protocol down event triggers the action specified under the **action** statement.

## 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 Release 8.5.

## RELATED DOCUMENTATION

| [Configuring an OAM Action Profile](#)

## premium (Output Priority Map)

### IN THIS SECTION

- [Syntax | 1093](#)
- [Hierarchy Level | 1093](#)

- Description | 1093
- Required Privilege Level | 1093
- Release Information | 1093

## Syntax

```
premium {
 forwarding-class class-name {
 loss-priority (high | low);
 }
}
```

## Hierarchy Level

```
[edit interfaces interface-name together-options ethernet-switch-profile
ethernet-policer-profile output-priority-map classifier]
```

## Description

For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[input-priority-map | 920](#)

## premium (Policer)

### IN THIS SECTION

- [Syntax | 1094](#)
- [Hierarchy Level | 1094](#)
- [Description | 1094](#)
- [Required Privilege Level | 1095](#)
- [Release Information | 1095](#)

### Syntax

```
premium {
 bandwidth-limit bps;
 burst-size-limit bytes;
}
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile
ethernet-policer-profile policer cos-policer-name]
```

### Description

Define a policer to apply to nonpremium traffic.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 271](#)

[aggregate \(Gigabit Ethernet CoS Policer\) | 798](#)

[ieee802.1p | 908](#)

## proxy

### IN THIS SECTION

- [Syntax | 1095](#)
- [Hierarchy Level | 1096](#)
- [Description | 1096](#)
- [Options | 1096](#)
- [Required Privilege Level | 1096](#)
- [Release Information | 1096](#)

## Syntax

```
proxy inet-address address;
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family tcc],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family family tcc]
```

## Description

For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 router.

Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. See [Configuring VLAN TCC Encapsulation](#) for more details.

## Options

**inet-address**—Configure the IP address of the neighbor to the TCC router.

## Required Privilege Level

**interface**—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

### Release History Table

| Release | Description                                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20.1R1  | Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. |

## RELATED DOCUMENTATION

[Configuring TCC Interface Switching](#)

[Configuring TCC Interface Switching](#)

[remote](#)

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

## rebalance (Aggregated Ethernet Interfaces)

### IN THIS SECTION

- [Syntax | 1097](#)
- [Hierarchy Level | 1097](#)
- [Description | 1097](#)
- [Options | 1098](#)
- [Required Privilege Level | 1098](#)
- [Release Information | 1098](#)

### Syntax

```
rebalance interval
```

### Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-
stateful per-flow]
```

### Description

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

## Options

*interval* Number of minutes after which the load-balancing state must be cleared for the specified interface.

- **Range:** 1 through 1000 flows per minute

## 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 Release 13.2R1.

## RELATED DOCUMENTATION

| *Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces*

## receive-options-packets

### IN THIS SECTION

- [Syntax | 1099](#)
- [Hierarchy Level | 1099](#)
- [Description | 1099](#)
- [Required Privilege Level | 1099](#)
- [Release Information | 1099](#)

## Syntax

```
receive-options-packets;
```

## Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family family inet]
```

## Description

For a Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Enabling Passive Monitoring on ATM Interfaces](#)

[Enabling Packet Flow Monitoring on SONET/SDH Interfaces](#)

## receive-ttl-exceeded

### IN THIS SECTION

- [Syntax | 1100](#)
- [Hierarchy Level | 1100](#)
- [Description | 1100](#)
- [Required Privilege Level | 1100](#)
- [Release Information | 1101](#)

### Syntax

```
receive-ttl-exceeded;
```

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family family inet]
```

### Description

For Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

[Enabling Passive Monitoring on ATM Interfaces](#)

[Enabling Packet Flow Monitoring on SONET/SDH Interfaces](#)

## recovery

### IN THIS SECTION

- [Syntax | 1101](#)
- [Hierarchy Level | 1101](#)
- [Description | 1102](#)
- [Options | 1102](#)
- [Required Privilege Level | 1102](#)
- [Release Information | 1103](#)

## Syntax

```
recovery {
 (auto | manual);
 timer timer-value;
}
```

## Hierarchy Level

```
[edit interfaces interfaces-name link-degrade-monitor]
```

## Description

Configure the mechanism to be used to recover a degraded link. The recovery options supported are auto and manual.

## Options

**auto** Recover a degraded link automatically. Use this option with the media-based action when there are no Layer 2 or Layer 3 protocols configured on the interface. If this option is configured, the degraded link is monitored at user-configured intervals; and if the link quality is found to have improved (if bit error rate hits the clear threshold), the link is automatically recovered. With this configuration, you must configure a timer value.

**manual** Recover a degraded link manually. Use this option with the media-based action configuration when Layer 2 and Layer 3 protocols are configured on the interface. If this option is configured, you need to use the **request interface link-degrade-recover *interface-name*** statement to recover the link.

**NOTE:** The manual recovery option is recommended for user deployments that have static route configurations causing the remote end of the link to start forwarding packets (as soon as the physical link is up) while autorecovery is in progress.

**timer**  
***timer-***  
***value*** Specify the interval value (in seconds) after which autorecovery of the degraded link must be triggered. This option is applicable if you configure the autorecovery option. The interval period starts from the time the link is degraded. The default interval is 1800 seconds. The autorecovery attempt is repeated until the link is recovered or the link monitoring feature is disabled through configuration.

**NOTE:** During autorecovery, you might notice link flaps at the remote end of the link.

## 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 Release 15.1.

### RELATED DOCUMENTATION

[Link Degradation Monitoring Overview | 600](#)

[link-degrade-monitor | 955](#)

[thresholds | 1179](#)

[request interface link-degrade-recover | 1243](#)

## remote-loopback

### IN THIS SECTION

- [Syntax | 1103](#)
- [Hierarchy Level | 1103](#)
- [Description | 1104](#)
- [Required Privilege Level | 1104](#)
- [Release Information | 1104](#)

### Syntax

```
remote-loopback;
```

### Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name]
```

## Description

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, set the remote DTE into loopback mode. Remove the statement from the configuration to take the remote DTE out of loopback mode. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

## 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 Release 8.2.

## RELATED DOCUMENTATION

| *Setting a Remote Interface into Loopback Mode*

## restore-interval

### IN THIS SECTION

- [Syntax | 1105](#)
- [Hierarchy Level | 1105](#)
- [Description | 1105](#)
- [Options | 1105](#)
- [Required Privilege Level | 1105](#)
- [Release Information | 1105](#)

## Syntax

```
restore-interval number;
```

## Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

## Description

Configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs).. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

## Options

*number*—Specify the restore interval.

- **Range:** 1 through 12 minutes

## 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 Release 9.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

## RELATED DOCUMENTATION

[Ethernet Ring Protection Switching Overview](#)

[Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)

[Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS](#)

[Configuring Ethernet Ring Protection Switching on Switches \(CLI Procedure\)](#)

## revertive

### IN THIS SECTION

- [Syntax | 1106](#)
- [Hierarchy Level | 1106](#)
- [Description | 1106](#)
- [Required Privilege Level | 1107](#)
- [Release Information | 1107](#)

### Syntax

```
revertive;
```

### Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

### Description

Enable the ability to switch to a better priority link (if one is available).

**NOTE:** By default, LACP link protection is revertive. However, you can use this statement to define a specific aggregated Ethernet interface as revertive to override a global **non-revertive** statement specified at the **[edit chassis]** hierarchy level.

## 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 Release 9.3.

## RELATED DOCUMENTATION

[non-revertive \(Chassis\) | 1037](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## routing-instance

### IN THIS SECTION

- [Syntax | 1108](#)
- [Hierarchy Level | 1108](#)
- [Description | 1108](#)
- [Default | 1108](#)
- [Required Privilege Level | 1108](#)
- [Release Information | 1108](#)

## Syntax

```
routing-instance {
 destination routing-instance-name;
}
```

## Hierarchy Level

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

## Description

To configure **interfaces** and **logical-systems**, specify the destination routing instance that points to the routing table containing the tunnel destination address.

## Default

The default Internet routing table is **inet.0**.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Junos OS Services Interfaces Library for Routing Devices](#)

## routing-instance (PPPoE Service Name Tables)

### IN THIS SECTION

- [Syntax | 1109](#)
- [Hierarchy Level | 1109](#)
- [Description | 1109](#)
- [Options | 1110](#)
- [Required Privilege Level | 1110](#)
- [Release Information | 1110](#)

### Syntax

```
routing-instance routing-instance-name;
```

### Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-
specifier aci circuit-id-string ari remote-id-string]
```

### Description

Use in conjunction with the **dynamic-profile** statement at the same hierarchy levels to specify the routing instance in which to instantiate a dynamic PPPoE interface. You can associate a routing instance with a named service entry, **empty** service entry, or **any** service entry configured in a PPPoE service name table, or with an agent circuit identifier/agent remote identifier (ACI/ARI) pair defined for these services.

The routing instance associated with a service entry in a PPPoE service name table overrides the routing instance associated with the PPPoE underlying interface on which the dynamic PPPoE interface is created.

If you include the **routing-instance** statement at the `[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci circuit-id-string ari remote-id-string]` hierarchy level, you

cannot also include the **static-interface** statement at this level. The **routing-instance** and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

## Options

***routing-instance-name***—Name of the routing instance in which the router instantiates the dynamic PPPoE interface.

## 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 Release 10.2.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

*Assigning a Dynamic Profile and Routing Instance to a Service Name or ACI/ARI Pair for Dynamic PPPoE Interface Creation*

## rx-enable

### IN THIS SECTION

- [Syntax | 1111](#)
- [Hierarchy Level | 1111](#)
- [Description | 1111](#)
- [Default | 1111](#)
- [Required Privilege Level | 1111](#)
- [Release Information | 1111](#)

## Syntax

```
expected-defect {
 rx-enable ;
}
```

## Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

## Description

Enable the ethernet expected defect (ETH-ED) function to process the received EDM PDUs.

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

## Default

The MEP does not process EDM PDUs.

## 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 Release 19.1.

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | 830

---

*show oam ethernet connectivity-fault-management mep-database*

## rx-max-duration

### IN THIS SECTION

- [Syntax | 1112](#)
- [Hierarchy Level | 1112](#)
- [Description | 1112](#)
- [Options | 1112](#)
- [Required Privilege Level | 1113](#)
- [Release Information | 1113](#)

### Syntax

```
expected-defect {
 rx-max-duration ;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

### Description

Duration to indicate the maximum acceptable value at which the loss of continuity alarms are suppressed. If the duration in the received EDM PDU exceeds this configured value then the duration value will be truncated to this configured value and loss of continuity (LoC) alarms shall be suppressed for this duration.

### Options

**Minimum value** The minimum value at which the loss of continuity alarms will be suppressed is 120 *seconds*.

**Minimum value** The maximum acceptable value at which the loss of continuity alarms will be suppressed is 3600 *seconds*.

**Default** 900 *seconds*.

### 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 Release 19.1.

### RELATED DOCUMENTATION

[connectivity-fault-management](#) | 830

*show oam ethernet connectivity-fault-management mep-database*

## sa-multicast (100-Gigabit Ethernet)

### IN THIS SECTION

- [Syntax](#) | 1113
- [Hierarchy Level](#) | 1114
- [Description](#) | 1114
- [Required Privilege Level](#) | 1114
- [Release Information](#) | 1114

### Syntax

```
sa-multicast;
```

## Hierarchy Level

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

## Description

Configure the 100-Gigabit Ethernet PIC or MIC to interoperate with other Juniper Networks 100-Gigabit Ethernet PICs.

**NOTE:** 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.

sa-multicast supports interoperability between the following PICs and MICs:

- 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP) and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) .
- 100-Gigabit Ethernet MICs and the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4).

## 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 Release 10.4.

## RELATED DOCUMENTATION

[Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | 238](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 239](#)

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

[Interoperability Between MPC4E \(MPC4E-3D-2CGE-8XGE\) and 100-Gigabit Ethernet PICs on Type 4 FPC](#)

[Configuring MPC4E \(MPC4E-3D-2CGE-8XGE\) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode](#)

[Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | 241](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 242](#)

[forwarding-mode \(100-Gigabit Ethernet\) | 886](#)

[sa-multicast \(PTX Series Packet Transport Routers\) | 1115](#)

[vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1226](#)

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 231](#)

## sa-multicast (PTX Series Packet Transport Routers)

### IN THIS SECTION

- [Syntax | 1115](#)
- [Hierarchy Level | 1115](#)
- [Description | 1116](#)
- [Required Privilege Level | 1116](#)
- [Release Information | 1116](#)

### Syntax

```
sa-multicast;
```

### Hierarchy Level

```
[edit chassis fpc slot pic slot port port-number forwarding-mode]
```

## Description

Configure source address (SA) multicast bit mode on the 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP to enable interoperability with 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4.

**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.

## 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 Release 12.1X48R4.

## RELATED DOCUMENTATION

[Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | 241](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 242](#)

## send-critical-event

### IN THIS SECTION

- [Syntax | 1117](#)
- [Hierarchy Level | 1117](#)
- [Description | 1117](#)
- [Required Privilege Level | 1117](#)
- [Release Information | 1117](#)

## Syntax

```
send-critical-event;
```

## Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile action]
```

## Description

Send OAM PDUs with the critical event bit set.

## 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 Release 8.5.

## RELATED DOCUMENTATION

| *Specifying the Actions to Be Taken for Link-Fault Management Events*

## server

### IN THIS SECTION

- [Syntax | 1118](#)
- [Hierarchy Level | 1118](#)
- [Description | 1118](#)

- Required Privilege Level | 1118
- Release Information | 1118

## Syntax

```
server;
```

## Hierarchy Level

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

## Description

Configure the router to operate in the PPPoE server mode. Supported on M120 and M320 Multiservice Edge Routers and MX Series 5G Universal Routing Platforms operating as access concentrators.

## 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 Release 8.5.

## RELATED DOCUMENTATION

| [Configuring the PPPoE Server Mode](#)

## service (PPPoE)

### IN THIS SECTION

- [Syntax | 1119](#)
- [Hierarchy Level | 1120](#)
- [Description | 1120](#)
- [Default | 1120](#)
- [Options | 1120](#)
- [Required Privilege Level | 1120](#)
- [Release Information | 1121](#)

### Syntax

```
service service-name {
 drop;
 delay seconds;
 terminate;
 dynamic-profile profile-name;
 routing-instance routing-instance-name;
 max-sessions number;
 agent-specifier {
 aci circuit-id-string ari remote-id-string {
 drop;
 delay seconds;
 terminate;
 dynamic-profile profile-name;
 routing-instance routing-instance-name;
 static-interface interface-name;
 }
 }
}
```

## Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name]
```

## Description

Specify the action taken by the interface on receipt of a PPPoE Active Discovery Initiation (PADI) control packet for the specified named service, **empty** service, or **any** service in a PPPoE service name table. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service.

## Default

The default action is terminate.

## Options

*service-name*—Service entry in the PPPoE service name table:

- *service-name*—Named service entry of up to 32 characters; for example, **premiumService**. You can configure a maximum of 512 named service entries across all PPPoE service name tables on the router.
- **empty**—Service entry of zero length that represents an unspecified service. Each PPPoE service name table includes one **empty** service entry by default.
- **any**—Default service for non-empty service entries that do not match the named or **empty** service entries configured in the PPPoE service name table. Each PPPoE service name table includes one **any** service entry by default.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 10.0.

**any**, **dynamic-profile**, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

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

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

*Configuring the Action Taken for the Any Service*

## service-name

### IN THIS SECTION

- [Syntax | 1121](#)
- [Hierarchy Level | 1122](#)
- [Description | 1122](#)
- [Options | 1122](#)
- [Required Privilege Level | 1122](#)
- [Release Information | 1122](#)

## Syntax

```
service-name name;
```

## Hierarchy Level

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

## Description

PPP over Ethernet interfaces, configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.

## Options

***name***—Service to be requested from the PPP over Ethernet server.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Configuring PPPoE | 48](#)

## service-name-table

### IN THIS SECTION

● [Syntax | 1123](#)

- Hierarchy Level | 1123
- Description | 1123
- Options | 1124
- Required Privilege Level | 1124
- Release Information | 1124

## Syntax

```
service-name-table table-name;
```

## Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number
family pppoe],
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-
number family pppoe],
[edit interfaces interface-name unit logical-unit-number family pppoe],
[edit interfaces interface-name unit logical-unit-number pppoe-underlying-
options],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family pppoe],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number pppoe-underlying-options]
```

## Description

Specify the PPPoE service name table assigned to a PPPoE underlying interface. This underlying interface is configured with either the **encapsulation ppp-over-ether** statement or the **family pppoe** statement; the two statements are mutually exclusive.

**NOTE:** The [edit ... family pppoe] hierarchies are supported only on MX Series routers with MPCs.

## Options

*table-name*—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

## 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 Release 10.0.

Support at the `[edit ... family pppoe]` hierarchies introduced in Junos OS Release 11.2.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

*Assigning a Service Name Table to a PPPoE Underlying Interface*

*Configuring the PPPoE Family for an Underlying Interface*

## service-name-tables

### IN THIS SECTION

- [Syntax | 1125](#)
- [Hierarchy Level | 1125](#)
- [Description | 1125](#)
- [Options | 1125](#)
- [Required Privilege Level | 1126](#)
- [Release Information | 1126](#)

## Syntax

```

service-name-tables table-name {
 service service-name {
 drop;
 delay seconds;
 terminate;
 dynamic-profile profile-name;
 routing-instance routing-instance-name;
 max-sessions number;
 agent-specifier {
 aci circuit-id-string ari remote-id-string {
 drop;
 delay seconds;
 terminate;
 dynamic-profile profile-name;
 routing-instance routing-instance-name;
 static-interface interface-name;
 }
 }
 }
}

```

## Hierarchy Level

```
[edit protocols pppoe]
```

## Description

Create and configure a PPPoE service name table. Specify the action taken for each service and remote access concentrator on receipt of a PPPoE Active Discovery Initiation (PADI) packet. You can also specify the dynamic profile and routing instance that the router uses to instantiate a dynamic PPPoE interface, and the maximum number of active PPPoE sessions that the router can establish with the specified service. A maximum of 32 PPPoE service name tables is supported per router.

## Options

**table-name**—Name of the PPPoE service name table, a string of up to 32 alphanumeric characters.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 10.0.

**dynamic-profile**, **routing-instance**, **max-sessions**, and **static-interface** options introduced in Junos OS Release 10.2.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

*Creating a Service Name Table*

## session-expiry (MX Series in Enhanced LAN Mode)

### IN THIS SECTION

- [Syntax | 1127](#)
- [Hierarchy Level | 1127](#)
- [Description | 1127](#)
- [Options | 1127](#)
- [Required Privilege Level | 1127](#)
- [Release Information | 1127](#)

## Syntax

```
session-expiry seconds;
```

## Hierarchy Level

```
[edit protocols authentication-access-control interface (all | [interface-
names])]]
```

## Description

Configure the maximum duration in seconds of a session.

## Options

*seconds*—Duration of session.

- **Range:** 1 through 65535
- **Default:** 3600

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 14.2.

## sonet

### IN THIS SECTION

- [Syntax | 1128](#)
- [Hierarchy Level | 1128](#)
- [Description | 1128](#)
- [Options | 1128](#)
- [Required Privilege Level | 1128](#)
- [Release Information | 1129](#)

### Syntax

```
sonet {
 device-count number;
}
```

### Hierarchy Level

```
[edit chassis aggregated-devices]
```

### Description

Configure properties for SONET/SDH aggregated devices on the router.

### Options

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices](#)

## source-address-filter

### IN THIS SECTION

- [Syntax | 1129](#)
- [Hierarchy Level | 1129](#)
- [Description | 1130](#)
- [Options | 1130](#)
- [Required Privilege Level | 1130](#)
- [Release Information | 1130](#)

## Syntax

```
source-address-filter {
 mac-address;
}
```

## Hierarchy Level

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

## Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the **source-filtering** statement in the configuration to enable source address filtering.

## Options

**mac-address**—MAC address filter. You can specify the MAC address as *nn.nn.nn.nn.nn* or *nnnn.nnnn.nnnn*, where *n* is a decimal digit. To specify more than one address, include multiple **mac-address** options in the **source-address-filter** statement.

If you enable the VRRP on a Fast Ethernet or Gigabit Ethernet interface, as described in [VRRP and VRRP for IPv6 Overview](#), and if you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from **00:00:5e:00:01:00** through **00:00:5e:00:01:ff** are reserved for VRRP, as defined in RFC 3768, *Virtual Router Redundancy Protocol*. When you configure the VRRP group, the group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

---

[Configuring MAC Address Filtering for Ethernet Interfaces | 20](#)

---

[Configuring MAC Address Filtering on PTX Series Packet Transport Routers | 22](#)

---

## source-filtering

### IN THIS SECTION

- [Syntax](#) | 1131
- [Hierarchy Level](#) | 1131
- [Description](#) | 1131
- [Default](#) | 1132
- [Required Privilege Level](#) | 1132
- [Release Information](#) | 1132

### Syntax

```
(source-filtering | no-source-filtering);
```

### Hierarchy Level

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

### Description

For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces only, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the **source-address-filter** statement.

If the remote Ethernet card is changed, the interface is no longer able to receive packets from the new card because it has a different MAC address.

## Default

Source address filtering is disabled.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

---

[Configuring MAC Address Filtering for Ethernet Interfaces | 20](#)

---

[Configuring MAC Address Filtering on PTX Series Packet Transport Routers | 22](#)

---

[accept-source-mac | 784](#)

---

[source-address-filter | 1129](#)

## speed

### IN THIS SECTION

- [Syntax | 1133](#)
- [Hierarchy Level | 1133](#)
- [Description | 1133](#)
- [Options | 1133](#)
- [Required Privilege Level | 1133](#)
- [Release Information | 1133](#)

## Syntax

```
speed [10G | 40G | 100G];
```

## Hierarchy Level

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

## Description

Configure the port speed on interface modules that support multiple port speeds. To check the port speed, use the **show interfaces** command. To determine whether a PIC has specific port speed configuration requirements, see the PIC's description in [PTX Series Interface Module Reference](#).

## Options

|      |          |
|------|----------|
| 10G  | 10 Gbps  |
| 40G  | 40 Gbps  |
| 100G | 100 Gbps |

**NOTE:** For PTX 1000 routers, the default port speed is 10 Gbps.

## 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 Release 15.1F3 and 16.1R2.

Statement introduced in Junos OS Release 16.1X65 for PTX1000 routers.

## RELATED DOCUMENTATION

[Modes of Operation | 204](#)

[Configuring Mixed-Rate Mode Operation | 205](#)

[mixed-rate-mode | 1014](#)

## speed

### IN THIS SECTION

- [Syntax | 1134](#)
- [Hierarchy Level | 1134](#)
- [Description | 1134](#)
- [Default | 1135](#)
- [Options | 1135](#)
- [Required Privilege Level | 1136](#)
- [Release Information | 1137](#)

### Syntax

```
speed (oc3-stm1 | oc12-stm4 | oc48-stm16 | 10G | 25G | 40G | 50G | 100G | 200G |
400G) ;
```

### Hierarchy Level

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

### Description

Configure the port speed for the ports on a line card or a device with a built-in FPC.

## Default

**oc3-stm1**

## Options

|                   |                                                 |
|-------------------|-------------------------------------------------|
| <b>100g</b>       | Supported ports operate at 100-Gbps speed.      |
| <b>10g</b>        | Supported ports operate at 10-Gbps speed.       |
| <b>1g</b>         | Supported ports operate at 1-Gbps speed.        |
| <b>200g</b>       | Supported ports operate at 200-Gbps speed.      |
| <b>25g</b>        | Supported ports operate at 25-Gbps speed.       |
| <b>400g</b>       | Supported ports operate at 400-Gbps speed.      |
| <b>40g</b>        | Supported ports operate at 40-Gbps speed.       |
| <b>50g</b>        | Supported ports operate at 50-Gbps speed.       |
| <b>oc12-stm4</b>  | Supported ports operate at OC12 or STM4 speed.  |
| <b>oc3-stm1</b>   | Supported ports operate at OC3 or STM1 speed.   |
| <b>oc48-stm16</b> | Supported ports operate at OC48 or STM16 speed. |

**NOTE:** You can configure the **oc12-stm4**, **oc3-stm1**, and **oc48-stm16** port speed options for SONET/SDH OC3/STM1 (Multi-Rate) MICs. However, for Channelized SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP and ATM MICs, you can configure only the **oc12-stm4** and **oc3-stm1** port speed options.

(MX Series with MPCs and ATM MICs with SFP) To configure up to OC12 CBR bandwidth speed per virtual circuit (VC) on an ATM MIC with SFP (MIC-3D-8OC3-2OC12-ATM), specify **oc12-stm4** as the speed for the specified port. You can configure the **oc12-stm4** port speed option only for ports 0 and 4 on an ATM MIC. If you configure the **oc12-stm4** port speed option for port 0, then ports 1, 2, and 3 are disabled. Similarly, if you configure the **oc12-stm4** port speed for port 4, then ports 5, 6, and 7 are disabled.

(MX Series with MPC7E-MRATE) To configure 100-Gbps, 10-Gbps, or 40-Gbps speed per port on an MPC7E (Multi-Rate) MPC, specify **100G**, **10G**, or **40G**, respectively, as the speed for the specified port. You can configure **10G** and **40G** port speed options on all the six ports of PIC 0 and PIC 1 of an MPC7E-

MRATE MPC. However, you can configure the **100G** port speed option only for ports 2 and 5 of PIC 0 and PIC 1 of an MPC7E-MRATE MPC.

(MX10003 routers with MX10003 MPC) To configure 100 Gbps, 10 Gbps, and 40 Gbps speed on all supported ports, specify 100G, 10G, or 40G, respectively, as the speed for the specified PIC. All the six ports of the fixed port PIC support 10-Gbps and 40-Gbps speeds. All the 12 ports of the Multi-rate MIC support 100-Gbps, 10-Gbps and 40-Gbps speeds. For more information see [MX10003 MPC \(Multi-Rate\)](#) and [Supported Active Physical Ports for Configuring Port Speed to Prevent Oversubscription on MX10003 MPC](#).

(MX204 Routers) The MX204 has four rate-selectable ports (referred to a PIC 0 ports) that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet port, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The MX204 also has eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports). For more information, see [MX204 Router Overview](#) and [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router](#)

(PTX1000 routers) The PTX1000 routers has 72 network ports that support the QSFP+ transceivers and can be configured as channelized 4x10-Gigabit Ethernet interfaces by default (for a maximum of 288 10-Gigabit Ethernet ports). When you configure the speed of a port, you must reset the FPC for the configuration to take effect. If you do not reset the FPC after modifying the port speed configuration, the router displays the following error message:

**pic\_need\_skip\_port\_speed\_change: pic-0/0: Port speed changed while previous port speed transition is pending. This PIC needs to be restarted. Please take this PIC or FPC offline and online via CLI.**

(PTX10001-36MR routers) See "[Port Speed on PTX10001-36MR Router Overview](#)" on page 350 to learn about multiple port speeds supported on PTX10001-36MR router, guidelines, and how to configure the port speed.

(PTX10003-80C and PTX10003-160C) To configure 100-Gbps, 10-Gbps, or 40-Gbps speed per port on the PTX10003-80C and PTX10003-160C line cards, specify **100g**, **10g**, or **40g**, respectively, as the speed for the specified port. Use the **number-of-sub-ports** statement to configure the number of channels for a particular port if the optics are used in a channelized mode. Use the **number-of-sub-ports** statement with the **speed** statement to configure a supported for the different channels based on the optics used. See "[PTX10003 Router Port Speed Overview](#)" on page 370 for more details.

(ACX5448-D) To configure 100-Gbps, 10-Gbps, or 40-Gbps speed per port on the ACX5448-D router, specify **100g**, **10g**, or **40g**, respectively, as the speed for the specified port.

## 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 Release 11.2.

Support for MPC7E (Multi-Rate) introduced in Junos OS Release 15.1F4.

Support for MX10003 MPC introduced in Junos OS Release 17.3R1.

Support for MX204 routers introduced in Junos OS Release 17.4R1.

Speed option 10Gbps, 40Gbps, and 100Gbps introduced in Junos OS Evolved Release 19.1R1 for PTX10003-80C, PTX10003-160C routers.

Speed options 100-Gbps, 40-Gbps, 25-Gbps, and 10-Gbps introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

## RELATED DOCUMENTATION

[Configuring Port Speed on Multi-Rate MICs](#)

[Configuring Port Speed on Routing Devices and Line Cards](#) | 418

[Configuring Port Speed on Routing Devices and Line Cards](#) | 418

[Introduction to Port Speed](#) | 370

[speed \(Gigabit Ethernet interface\)](#) | 1150

## speed (Ethernet)

### IN THIS SECTION

- [Syntax \(ACX Series, EX Series, MX Series\)](#) | 1138
- [Syntax \(ACX5448, ACX710\)](#) | 1138
- [Syntax \(ACX5448-D\)](#) | 1138
- [Syntax \(EX Series\)](#) | 1138
- [Syntax \(EX2300\)](#) | 1138
- [Syntax \(EX4300\)](#) | 1139
- [Syntax \(EX4600, OCX1100, QFX Series\)](#) | 1139
- [Syntax \(QFX5100-48T\)](#) | 1139

- Syntax (QFX5130-32CD) | **1139**
- Hierarchy Level (ACX Series, EX Series, MX Series) | **1139**
- Hierarchy Level (EX Series) | **1139**
- Hierarchy Level (ACX5448, ACX5448-D, ACX710, EX2300, EX4300, EX4600, OCX Series, QFX Series, QFX5100-48T, QFX5130-32CD) | **1139**
- Description | **1140**
- Options | **1140**
- Required Privilege Level | **1147**
- Release Information | **1147**

### Syntax (ACX Series, EX Series, MX Series)

```
speed (10m | 100m | 10g | 1g | 2.5g | 5g | auto | auto-10m-100m);
```

### Syntax (ACX5448, ACX710)

```
speed (10g | 25g | 40g | 100g | 100m | 1g | auto);
```

### Syntax (ACX5448-D)

```
speed (10g | 25g | 40g | 100g | 1g | auto);
```

### Syntax (EX Series)

```
speed (auto-negotiation | speed) ;
```

### Syntax (EX2300)

```
speed (10m | 100m | 1g | 2.5g);
```

**Syntax (EX4300)**

```
speed (10m | 100m | 1g | 2.5g | 5g | 10g);
```

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

```
speed (10g | 1g | 100m);
```

**Syntax (QFX5100-48T)**

```
speed (10g | 1g | 100m | auto);
```

**Syntax (QFX5130-32CD)**

```
speed (10g | 25g | 40g | 100g | 400g);
```

**Hierarchy Level (ACX Series, EX Series, MX Series)**

```
[edit interfaces interface-name],
[edit interfaces ge-pim/0/0 switch-options switch-port port-number]
```

**Hierarchy Level (EX Series)**

```
[edit interfaces interface-name ether-options]
```

**Hierarchy Level (ACX5448, ACX5448-D, ACX710, EX2300, EX4300, EX4600, OCX Series, QFX Series, QFX5100-48T, QFX5130-32CD)**

```
[edit interfaces interface-name]
```

## Description

Configure the interface speed. This statement applies to the following interfaces:

- Management Ethernet interface (**fxp0** or **em0**)
- Fast Ethernet 12-port and 48-port PICs
- Built-in Fast Ethernet port on the FIC (M7i router)
- Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers
- Gigabit Ethernet interfaces on EX Series switches

If you enable autonegotiation, then the device automatically negotiates the speed based on the speed of the other end of the link. [Table 132 on page 1140](#) describes the autonegotiation option available for different platforms:

**Table 132: Autonegotiation Options**

| Autonegotiation Option  | Description                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>auto-negotiation</b> | Automatically negotiate the speed based on the speed of the other end of the link.                                |
| <b>auto</b>             | Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link. |
| <b>auto-10m-100m</b>    | Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link.          |

See *Speed and Autonegotiation* for more details.

## Options

Routers support autonegotiation by default. To enable or disable autonegotiation, see "[autonegotiation \(Routers\)](#)" on page 806. [Table 133 on page 1141](#) summarizes the speed and autonegotiation supported on different routing platforms:

Table 133: Speed and Autonegotiation support on Routers

| Router          | Speed                                                                                                                                                                       | Autonegotiation Supported (Yes/No) ? |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| ACX Series      | <ul style="list-style-type: none"> <li>• 10M</li> <li>• 100M</li> <li>• 1G</li> <li>• 2.5G</li> <li>• 5G</li> <li>• 10G</li> <li>• auto</li> <li>• auto-10m-100m</li> </ul> | Yes                                  |
| ACX5448, ACX710 | <ul style="list-style-type: none"> <li>• 10G</li> <li>• 25G</li> <li>• 40G</li> <li>• 100G</li> <li>• 100M</li> <li>• 1G</li> <li>• auto</li> </ul>                         | Yes                                  |

Table 133: Speed and Autonegotiation support on Routers (Continued)

| Router    | Speed                                                                                                                                                                       | Autonegotiation Supported (Yes/No) ? |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| ACX5448-D | <ul style="list-style-type: none"> <li>• 10G</li> <li>• 25G</li> <li>• 40G</li> <li>• 100G</li> <li>• 1G</li> <li>• auto</li> </ul>                                         | Yes                                  |
| MX Series | <ul style="list-style-type: none"> <li>• 10M</li> <li>• 100M</li> <li>• 1G</li> <li>• 2.5G</li> <li>• 5G</li> <li>• 10G</li> <li>• auto</li> <li>• auto-10m-100m</li> </ul> | Yes                                  |

Switches support autonegotiation by default. To enable or disable autonegotiation, see *auto-negotiation (Switches)*. [Table 134 on page 1143](#) summarizes the speed and autonegotiation supported on different switching platforms:

Table 134: Speed and Autonegotiation support on Switches

| Switch                      | Speed                                                                                                                                                                                                                                                     | Autonegotiation Supported (Yes/No) ? |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| EX Series                   | <ul style="list-style-type: none"> <li>• 10M</li> <li>• 100M</li> <li>• 1G</li> <li>• 2.5G</li> <li>• 5G</li> <li>• 10G</li> <li>• auto</li> <li>• auto-10m-100m</li> </ul>                                                                               | Yes                                  |
| EX2300-48MP and EX2300-24MP | <ul style="list-style-type: none"> <li>• 10M – supported on EX series switches and only on <b>ge</b> interfaces of EX2300MP switch.</li> <li>• 100M</li> <li>• 1G</li> <li>• 2.5G – supported only on <b>mge</b> interfaces of E2300MP switch.</li> </ul> | Yes                                  |

Table 134: Speed and Autonegotiation support on Switches (*Continued*)

| Switch                       | Speed                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Autonegotiation Supported (Yes/No) ? |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| EX4300-48MP (EX-UM-4SFPP-MR) | <ul style="list-style-type: none"> <li>• 10M – supported only on <b>ge</b> interfaces.</li> <li>• 100M – supported on <b>ge</b> and <b>mge</b> interfaces.</li> <li>• 1G – supported on <b>ge</b>, <b>mge</b> interfaces, and 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches. The 1-Gbps speed is supported on the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module of EX4300-48MP switches from Junos OS Release 19.1R1 onwards.</li> <li>• 2.5G – supported only on <b>mge</b> interfaces.</li> <li>• 5G – supported only on <b>mge</b> interfaces.</li> <li>• 10G – supported on <b>mge</b> interfaces and 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches.</li> </ul> | Yes                                  |

Table 134: Speed and Autonegotiation support on Switches (*Continued*)

| Switch                                                  | Speed                                                                                                                                                                                                                                                                          | Autonegotiation Supported (Yes/No) ? |
|---------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| EX4600-40F (connected with JNP-SFPP-10GE-T transceiver) | <ul style="list-style-type: none"> <li>• 100M – supported only on <b>mge</b> interfaces.</li> <li>• 1G</li> <li>• 10G</li> </ul> <p>For more information, see <i>Autonegotiation support for EX4600-40F, QFX5110-48S and QFX5100-48S with JNP-SFPP-10GE-T transceiver.</i></p> | Yes                                  |
| OCX1100, QFX Series, QFabric                            | <ul style="list-style-type: none"> <li>• 100M</li> <li>• 1G</li> <li>• 10G</li> </ul>                                                                                                                                                                                          | No                                   |
| QFX5100-48T                                             | <ul style="list-style-type: none"> <li>• 100M</li> <li>• 1G</li> <li>• 10G</li> <li>• auto</li> </ul> <p>For more information, see <i>Configure Speed and Autonegotiation on QFX5100-48T Switches.</i></p>                                                                     | Yes                                  |

Table 134: Speed and Autonegotiation support on Switches *(Continued)*

| Switch                                                                   | Speed                                                                                                                                                                                                                               | Autonegotiation Supported (Yes/No) ? |
|--------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| QFX5110-48S (connected with QFX-SFP-1GE-T transceiver)                   | <ul style="list-style-type: none"> <li>• 100M</li> <li>• 1G</li> <li>• 10G</li> <li>• 40G</li> <li>• 100G</li> <li>• auto</li> </ul>                                                                                                | Yes                                  |
| QFX5110-48S and QFX5100-48S (connected with JNP-SFPP-10GE-T transceiver) | <ul style="list-style-type: none"> <li>• 100M</li> <li>• 1G</li> <li>• 10G</li> </ul> <p>For more information, see <i>Autonegotiation support for EX4600-40F, QFX5110-48S and QFX5100-48S with JNP-SFPP-10GE-T transceiver.</i></p> | Yes                                  |
| QFX5120-48Y (connected with JNP-SFPP-10GE-T transceiver)                 | <ul style="list-style-type: none"> <li>• 1G</li> <li>• 10G</li> </ul> <p>For more information, see <i>Autonegotiation support for QFX5120-48Y with JNP-SFPP-10GE-T transceiver.</i></p>                                             | No                                   |

**Table 134: Speed and Autonegotiation support on Switches (Continued)**

| Switch       | Speed                                                                                                                                                                                                          | Autonegotiation Supported (Yes/No) ? |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| QFX5130-32CD | <ul style="list-style-type: none"> <li>• 400G</li> <li>• 100G</li> <li>• 40G</li> <li>• 25G</li> <li>• 10G</li> </ul> <p>For more information, see <i>Channelizing Interfaces on QFX5200-32C Switches</i>.</p> | No                                   |

### Required Privilege Level

interface—To view this statement in the configuration.

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

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

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

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Speed option 2.5Gbps introduced in Junos OS Release 18.1R2 for EX2300 switch.

Speed option 10Gbps and 5Gbps introduced in Junos OS Release 18.2R1 for EX4300 switch.

Speed option 1-Gbps is introduced in Junos OS Release 19.1R1 on the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches.

Speed options 100-Mbps, 1-Gbps, and auto is introduced in Junos OS Releases 18.4R1S2, 18.4R2, and 19.2R1 and later for ACX5448 Universal Metro Routers.

Speed option 10Gbps, 40Gbps, and 100Gbps introduced in Junos OS Evolved Release 19.1R1 for PTX10003-80C, PTX10003-160C routers.

Speed options 100-Gbps, 40-Gbps, 25-Gbps, and 10-Gbps introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

## RELATED DOCUMENTATION

*Speed and Autonegotiation*

*Port Settings*

[Configuring the Interface Speed](#)

[Configuring the Interface Speed on Ethernet Interfaces | 6](#)

[Configuring Gigabit Ethernet Autonegotiation | 284](#)

*Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support auto-negotiation*

*Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches*

[Junos OS Network Interfaces Library for Routing Devices](#)

[Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

[Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\)](#)

[Junos OS Ethernet Interfaces Configuration Guide](#)

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

## speed (MX Series DPC)

### IN THIS SECTION

- [Syntax | 1149](#)
- [Hierarchy Level | 1149](#)
- [Description | 1149](#)
- [Options | 1149](#)
- [Required Privilege Level | 1149](#)

## Syntax

```
speed (auto | 1Gbps | 100Mbps | 10Mbps);
```

## Hierarchy Level

```
[edit interfaces ge-fpc/pic/port]
```

## Description

On MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs you can set auto negotiation of speed. To specify the auto negotiation speed, use the **speed (auto | 1Gbps | 100Mbps | 10Mbps)** statement under the **[edit interface ge-/fpc/pic/port]** hierarchy level. The **auto** option will attempt to automatically match the rate of the connected interface. To set port speed negotiation to a specific rate, set the port speed to **1Gbps**, **100Mbps**, or **10Mbps**.

**NOTE:** If the negotiated speed and the interface speed do not match, the link will not be brought up. Half duplex mode is not supported.

## Options

You can specify the speed as either **auto** (autonegotiate), **10Mbps** (10 Mbps), **100Mbps** (100 Mbps), or **1Gbps** (1 Gbps).

## 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 Release 9.5.

### RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Autonegotiation | 284](#)

[no-auto-mdix | 1030](#)

## speed (Gigabit Ethernet interface)

### IN THIS SECTION

- [Syntax | 1150](#)
- [Hierarchy Level | 1150](#)
- [Description | 1151](#)
- [Default | 1151](#)
- [Options | 1152](#)
- [Required Privilege Level | 1152](#)
- [Release Information | 1152](#)

### Syntax

```
speed (1G | 10G);
```

### Hierarchy Level

```
[edit interfaces intf-name gigether-options]
```

## Description

(MX10003, MX204, MX10008, and MX10016 routers) the 10-Gbps port can operate in 1-Gbps mode also. When a port operates in 10-Gbps speed, you can change the operating speed to 1-Gbps using the configuration **speed 1G** in this configuration statement. After you commit this configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed.

On fixed-port PIC and non-MACsec MIC of MX10003 router, you can configure one or all 10-Gbps port operating in 4X10-Gbps speed to operate in 1-Gbps speed. On MX204 routers, you can configure the 4X10-Gbps port on one of the fixed-port PICs to operate in 1-Gbps mode. And on the other fixed-port PIC, you can configure the 10-Gbps port to 1-Gbps speed. On MX10003 and MX204 routers, 1-Gbps speed is supported with **speed 1g** configuration. On MX204 routers, the 1-Gbps SFP port supports auto-negotiation. You can configure auto-negotiation by using the command **set interfaces interface-name gigether-options auto-negotiation**. For more information, see ["auto-negotiation" on page 806](#).

On JNP10K-LC2101 MPC, you can configure one or all 10-Gbps ports operating in 4x10-Gbps speed to operate in 1-Gbps speed. Also, autonegotiation is supported when the interface speed is configured for 1-Gbps speed.

### NOTE:

- On the MX10003 router, the MACsec MIC does not provide 1-Gbps speed. If you attempt to change the operating speed to 1-Gbps, syslog displays that this feature is not supported on the MACsec MIC.
- (MX10003, MX204, MX10008, and MX10016 routers) Rate selectability at PIC level and port level does not support 1-Gbps speed.
- (MX10003, MX204, MX10008, and MX10016 routers) The interface name prefix is **xe**.
- (MX10003, MX204, MX10008, and MX10016 routers) After configuring 1-Gbps speed, the protocol continues to advertise the bandwidth as 10-Gigabit Ethernet.
- On MX10003 and MX204 routers, Link Aggregation Group (LAG) is supported on 10-Gbps speed only. It is not supported on 1-Gbps speed.

To view the speed configured for the interface, execute the *show interfaces extensive* command. The **Speed Configuration** output parameter in the command output indicates the current operation speed of the interface. If the interface is configured with 1-Gbps speed, then **Speed Configuration** displays 1G; if the interface is configured with 10-Gbps speed, Speed Configuration displays AUTO.

## Default

10G

## Options

**1G**—Supported ports operate at 1-Gbps speed.

**10G**—Supported ports operate at 10-Gbps speed.

## 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 Release 18.1R1.

## RELATED DOCUMENTATION

[Configuring Port Speed on Multi-Rate MICs](#)

[Introduction to Port Speed | 370](#)

[Introduction to Port Speed | 370](#)

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

## speed (24-port and 12-port 10 Gigabit Ethernet PIC)

### IN THIS SECTION

- [Syntax | 1153](#)
- [Hierarchy Level | 1153](#)
- [Description | 1153](#)
- [Options | 1153](#)
- [Required Privilege Level | 1154](#)
- [Release Information | 1154](#)

## Syntax

speed 1G | 10G

## Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number]
[edit chassis fpc slot-number pic pic-number port port-number]
[edit chassis lcc number fpc slot-number pic pic-number mixed-rate-mode]
(Routing Matrix)
```

## Description

Configure the port speed on the following interface modules:

- PF-24XGE-SFPP or the PF-12XGE-SFPP PIC on a T4000 standalone router or on an LCC in a TX Matrix Plus routing matrix with 3D SIBs

**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.

- P1-PTX-24-10GE-SFPP PIC on the PTX3000 router
- P1-PTX-24-10GE-SFPP PIC with the FPC2-PTX-P1A on the PTX5000 router

Dual-rate support for the P1-PTX-24-10GE-SFPP enables you to switch all port speeds to either 1 Gbps or 10 Gbps. The default is 10 Gbps. All ports are configured to the same speed; there is no mixed-rate-mode capability. Changing the port speed causes the PIC to reboot.

To return all ports to the 10-Gbps port speed, use the **delete chassis fpc *fpc-number* pic *pic-number* speed 1G** statement. To check the port speed, use the **show interfaces** command.

**NOTE:** For the 1-Gbps port speed on the P1-PTX-24-10GE-SFPP PIC, you can use either the SFP-1GE-SX or the SFP-1GE-LX transceiver.

## Options

**1 G**

1 Gbps

10 G

10 Gbps

### 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 Release 13.3.

### RELATED DOCUMENTATION

[Modes of Operation | 204](#)

[Configuring Mixed-Rate Mode Operation | 205](#)

[mixed-rate-mode | 1014](#)

## static-interface

### IN THIS SECTION

- [Syntax | 1154](#)
- [Hierarchy Level | 1155](#)
- [Description | 1155](#)
- [Options | 1155](#)
- [Required Privilege Level | 1155](#)
- [Release Information | 1155](#)

### Syntax

```
static-interface interface-name;
```

## Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name agent-
specifier aci circuit-id-string ari remote-id-string]
```

## Description

Reserve the specified static PPPoE interface for use only by the PPPoE client with matching agent circuit identifier (ACI) and agent remote identifier (ARI) information. You can specify only one static interface per ACI/ARI pair configured for a named service entry, **empty** service entry, or **any** service entry in the PPPoE service name table.

The static interface associated with an ACI/ARI pair takes precedence over the general pool of static interfaces associated with the PPPoE underlying interface.

If you include the **static-interface** statement in the configuration, you cannot also include either the **dynamic-profile** statement or the **routing-instance** statement. The **dynamic-profile**, **routing-instance**, and **static-interface** statements are mutually exclusive for ACI/ARI pair configurations.

## Options

*interface-name*—Name of the static PPPoE interface reserved for use by the PPPoE client with matching ACI/ARI information. Specify the interface in the format **pp0.logical**, where *logical* is a logical unit number from 0 through 16385 for static interfaces.

## 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 Release 10.2.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

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

## switch-options

### IN THIS SECTION

- [Syntax | 1156](#)
- [Hierarchy Level | 1156](#)
- [Description | 1156](#)
- [Required Privilege Level | 1156](#)
- [Release Information | 1157](#)

### Syntax

```
switch-options {
 switch-port port-number {
 (auto-negotiation | no-auto-negotiation);
 speed (10m | 100m | 1g);
 link-mode (full-duplex | half-duplex);
 }
}
```

### Hierarchy Level

```
[edit interfaces ge-pim/0/0]
```

### Description

Configuration of the physical port characteristics is done under the single physical interface.

### 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 Release 8.4.

## switch-port

### IN THIS SECTION

- [Syntax | 1157](#)
- [Hierarchy Level | 1157](#)
- [Description | 1157](#)
- [Default | 1158](#)
- [Options | 1158](#)
- [Required Privilege Level | 1158](#)
- [Release Information | 1158](#)

## Syntax

```
switch-port port-number {
 (auto-negotiation | no-auto-negotiation);
 speed (10m | 100m | 1g);
 link-mode (full-duplex | half-duplex);
}
```

## Hierarchy Level

```
[edit interfaces ge-pim/0/0 switch-options]
```

## Description

Configuration of the physical port characteristics, done under the single physical interface.

## Default

Autonegotiation is enabled by default. If the link speed and duplex are also configured, the interfaces use the values configured as the desired values in the negotiation.

## Options

*port-number*—Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.4.

## symbol-period

### IN THIS SECTION

- [Syntax | 1159](#)
- [Hierarchy Level | 1159](#)
- [Description | 1159](#)
- [Options | 1159](#)
- [Required Privilege Level | 1159](#)
- [Release Information | 1159](#)

## Syntax

```
symbol-period count;
```

## Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event, link-
event-rate],
[edit protocols oam link-fault-management interface interface-name event-
thresholds]
```

## Description

Configure the threshold for sending symbol period events or taking the action specified in the action profile.

A symbol error is any symbol code error on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period window. The default period window is the number of symbols that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

## Options

*count*—Threshold count for symbol period events.

- **Range:** 0 through 100

## 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 Release 8.4.

## RELATED DOCUMENTATION

*Configuring Threshold Values for Local Fault Events on an Interface*

*Configuring Threshold Values for Fault Events in an Action Profile*

## symmetric-hash

### IN THIS SECTION

- [Syntax | 1160](#)
- [Hierarchy Level | 1160](#)
- [Description | 1160](#)
- [Options | 1161](#)
- [Required Privilege Level | 1161](#)
- [Release Information | 1161](#)

### Syntax

```
symmetric-hash {
 complement;
}
```

### Hierarchy Level

```
[edit chassis fpc slot-number pic slot-number hash-key family inet],
[edit chassis fpc slot-number pic slot-number hash-key family multiservice]
```

### Description

(MX Series 5G Universal Routing Platforms only) Configure the symmetric hash or symmetric hash complement at the PIC level for configuring symmetrical load balancing on an 802.3ad Link Aggregation Group.

## Options

**complement**      Include the complement of the symmetric hash in the hash key.

## 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 Release 9.6.

## RELATED DOCUMENTATION

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

## sync-reset

### IN THIS SECTION

- [Syntax | 1162](#)
- [Hierarchy Level | 1162](#)
- [Description | 1162](#)
- [Default | 1162](#)
- [Options | 1162](#)
- [Required Privilege Level | 1163](#)
- [Release Information | 1163](#)

## Syntax

```
sync-reset (disable | enable);
```

## Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options lacp],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-
ether-options lacp],
[edit interfaces name aggregated-ether-options lacp]
```

## Description

For redundant Ethernet interface link aggregation group links, you can configure the minimum number of physical child links on the primary node in the redundant Ethernet interface that must be working. If the minimum number of operating child links falls below the configured value, the interface is marked down even if some of the child interfaces are working. LACP marks these operating child interfaces or links that are working, as out of synchronization. This enables a peer switch, that does not have the minimum link configuration, to mark its interface as down as well. The peer switch can be a Juniper Switch or any other third party switch. As a result, the interface is marked as down on both the switches until the number of operating child links is more than the configured value.

For an aggregated ethernet interface, you cannot configure all three configuration options, **bfd-liveness-detection**, **minimum-links**, and **sync-reset** at the same time.

## Default

(ACX Series routers) Disabled.

(MX Series routers) Disabled.

(PTX Series routers) Disabled.

(QFX Series switches) Disabled.

(EX Series switches) Disabled.

## Options

**disable**                      Disable synchronization reset.

**enable** Enable synchronization reset.

## Required Privilege Level

interface

## Release Information

Statement introduced in Junos OS Release 16.1R1.

## RELATED DOCUMENTATION

[minimum-links](#) | [1012](#)

[show lacp interfaces](#) | [1756](#)

## syslog (OAM Action)

### IN THIS SECTION

- [Syntax](#) | [1163](#)
- [Hierarchy Level](#) | [1164](#)
- [Description](#) | [1164](#)
- [Required Privilege Level](#) | [1164](#)
- [Release Information](#) | [1164](#)

## Syntax

```
syslog;
```

## Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile action]
```

## Description

Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event.

Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.

## Required Privilege Level

interface—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 8.5.

## RELATED DOCUMENTATION

*Specifying the Actions to Be Taken for Link-Fault Management Events*

*Configuring Ethernet OAM Link Fault Management*

## system-id

### IN THIS SECTION

- [Syntax | 1165](#)
- [Hierarchy Level | 1165](#)

- Description | 1165
- Required Privilege Level | 1165
- Release Information | 1165

## Syntax

```
system-id system-id;
```

## Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp]
```

## Description

Define the LACP system identifier at the aggregated Ethernet interface level.

The user-defined system identifier in LACP enables two ports from two separate routers (M Series or MX Series routers) 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.

## 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 Release 12.2R1

## RELATED DOCUMENTATION

| [Configuring LACP for Aggregated Ethernet Interfaces](#)

## system-priority

### IN THIS SECTION

- [Syntax | 1166](#)
- [Hierarchy Level | 1166](#)
- [Description | 1166](#)
- [Options | 1167](#)
- [Required Privilege Level | 1167](#)
- [Release Information | 1167](#)

### Syntax

```
system-priority priority;
```

### Hierarchy Level

```
[edit chassis aggregated-devices ethernet lacp]
```

### Description

Define LACP system priority for aggregated Ethernet interfaces at the global (chassis) level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

## Options

*priority*—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

- **Range:** 0 through 65535
- **Default:** 127

## 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 Release 9.3.

## RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## system-priority

### IN THIS SECTION

- [Syntax | 1168](#)
- [Hierarchy Level | 1168](#)
- [Description | 1168](#)
- [Options | 1168](#)
- [Required Privilege Level | 1168](#)
- [Release Information | 1168](#)

## Syntax

```
system-priority priority;
```

## Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp]
```

## Description

Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global **[edit chassis]** hierarchy level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

## Options

*priority*—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

- **Range:** 0 through 65535
- **Default:** 127

## 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 Release 9.3.

## targeted-options (Grouping Subscribers by Bandwidth Usage)

### IN THIS SECTION

- [Syntax | 1169](#)
- [Hierarchy Level | 1169](#)
- [Description | 1169](#)
- [Options | 1170](#)
- [Required Privilege Level | 1170](#)
- [Release Information | 1170](#)

### Syntax

```
targeted-options {
 backup backup;
 group group;
 primary primary;
 weight ($junos-interface-target-weight | weight-value);
}
```

### Hierarchy Level

```
[edit dynamic-profiles name interfaces name unit logical-unit-number],
[edit dynamic-profiles name logical-systems name interfaces name unit logical-
unit-number],
[edit interfaces name unit logical-unit-number]
```

### Description

Configure primary and backup links, group similar subscribers, and specify a subscriber weight for manual targeting to distribute subscribers across aggregated Ethernet member links.

## Options

|                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>backup</b>                                                                     | (Optional) Specify a backup member link per subscriber when you configure manual targeting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>group</b>                                                                      | (Optional) Assign a group name for subscribers with similar bandwidth usage. Subscribers that are configured for targeted distribution without a group name are added to the <b>default</b> group and distributed evenly across member links. Grouping of subscribers is supported only for static subscribers. <ul style="list-style-type: none"> <li>• <b>Default:</b> default</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>primary</b>                                                                    | Specify a primary member link per subscriber when you configure manual targeting. You must always configure a primary link when you configure manual targeting.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>weight</b><br>( <b>\$junos-interface-target-weight</b>   <b>weight-value</b> ) | Specify the weight for targeted subscribers like PPPoe, demux, and conventional VLANs based on factors such as customer preferences, class of service (CoS), or bandwidth requirement. Member links for logical interfaces of aggregated Ethernet logical interfaces are assigned based on the value of the weight . When a new VLAN is added to the same aggregated Ethernet bundle, then the primary member link selected for targeting is the one with the minimum primary load and the backup link selected for targeting is the one with the minimum overall load. <p>The \$junos-interface-target-weight predefined variable is supported for dynamic configuration only. When you configure this predefined variable, the weight value is sourced from VSA 26-213 in the RADIUS Access-Accept message when a dynamic subscriber is authenticated.</p> <ul style="list-style-type: none"> <li>• <b>Range:</b> 1 through 1000</li> </ul> |

## 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 Release 16.1.

**weight** option added in Junos OS Release 17.3 for MX Series and MX Virtual Chassis.

**\$junos-interface-target-weight** option added in Junos OS Release 18.4R1.

## RELATED DOCUMENTATION

[Understanding Support for Targeted Distribution of Logical Interface Sets of Static VLANs over Aggregated Ethernet Logical Interfaces](#)

*Using RADIUS-Sourced Weights for Interface and Interface Set Targeted Distribution*

*RADIUS-Sourced Weights for Interface and Interface Set Targeted Distribution*

## targeted-options (Manual Targeting)

### IN THIS SECTION

- [Syntax | 1171](#)
- [Hierarchy Level | 1171](#)
- [Description | 1172](#)
- [Options | 1172](#)
- [Required Privilege Level | 1172](#)
- [Release Information | 1172](#)

### Syntax

```
targeted-options {
 (logical-interface-chassis-redundancy | logical-interface-fpc-redundancy);
 rebalance-periodic {
 interval interval;
 start-time start-time;
 }
 type (auto | manual);
}
```

### Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options],
[edit dynamic-profiles name logical-systems name interfaces name aggregated-
ether-options],
```

```
[edit interfaces name aggregated-ether-options]
[edit interfaces name unit logical-unit-number]
```

## Description

Configure manual targeting or auto-targeting.

## Options

**type** Configure manual targeting type as **manual** or **auto**.

- Values:
  - **auto**—Configure targeted-distribution without specific primary and backup links.
  - **manual**—Configure targeted distribution with specific member links as primary and backup for a subscriber. When you configure manual targeting, you must always configure a primary link. Configuring a backup link is optional. You specify the primary and backup links for a subscriber in the individual interface configuration. You configure primary and backup links by using the "[targeted-options](#)" on page 1169 statement at the `[edit interfaces name unit]` hierarchy level.

Manual targeting enhances the distribution of targeted VLANs or subscribers across member links of an aggregated Ethernet bundle by making it bandwidth-aware.

- **Default:** auto

The remaining statements are described separately.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 16.1.

## RELATED DOCUMENTATION

[Ethernet Interfaces Overview](#) | 2

[targeted-options \(Grouping Subscribers by Bandwidth Usage\) | 1169](#)

[Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis](#)

## targeted-distribution

### IN THIS SECTION

- [Syntax | 1173](#)
- [Hierarchy Level | 1173](#)
- [Description | 1173](#)
- [Options | 1174](#)
- [Required Privilege Level | 1174](#)
- [Release Information | 1174](#)

### Syntax

```
targeted-distribution primary-list primary-list [backup-list backup-list;
```

### Hierarchy Level

```
[edit logical-systems name interfaces name unit]
```

### Description

Configure egress data for a member link in an aggregated Ethernet bundle. Specify a distribution list as primary list and a different distribution list as backup list. A backup list is provisioned in the event the primary list goes down.

## Options

**primary-list** (Optional) Specify the role of the distribution list as primary. Member links of the aggregated Ethernet are assigned membership to the distribution list.

**backup-list** (Optional) Specify the role of the distribution list as backup. Member links of the aggregated Ethernet are assigned membership to the distribution list.

## 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 Release 16.1R1.

## RELATED DOCUMENTATION

[Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links](#)

[distribution-list](#)

[targeted-options](#) | [1174](#)

## targeted-options

### IN THIS SECTION

- [Syntax](#) | [1175](#)
- [Hierarchy Level](#) | [1175](#)
- [Description](#) | [1175](#)
- [Options](#) | [1175](#)
- [Required Privilege Level](#) | [1175](#)
- [Release Information](#) | [1175](#)

## Syntax

```
targeted-options {
 type (auto | manual);
}
```

## Hierarchy Level

```
[edit dynamic-profiles name interfaces name aggregated-ether-options]
[edit dynamic-profiles name logical-systems name interfaces name aggregated-
ether-options]
[edit interfaces name aggregated-ether-options]
```

## Description

Specify the type of targeting to be used for targeted distribution. Specify the targeting option as **manual** for conventional VLAN targeting. By default, the targeting option is **auto**.

## Options

**type** Specify the type of targeting to be used for targeted distribution.

- **Default: auto**—By default, targeted option is set to **auto**.
- Values:
  - **manual**—Use **manual** keyword to enforce manual targeting on conventional VLANs.

## 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 Release 16.1.

## RELATED DOCUMENTATION

[Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links](#)

[distribution-list](#)

[targeted-distribution](#) | 1173

## tdm-options (Interfaces)

### IN THIS SECTION

- [Syntax](#) | 1176
- [Hierarchy Level](#) | 1176
- [Description](#) | 1177
- [Options](#) | 1177
- [Required Privilege Level](#) | 1177
- [Release Information](#) | 1177

### Syntax

```
tdm-options {
 ces-psn-channel ...
 ces-psn-port-dmac-check-enable ;
 iwf-params...
 sfp-type sfptype;
 tdm-in-loop ;
 tdm-out-loop ;
}
```

### Hierarchy Level

```
[edit interfaces interface-name]
```

## Description

Configure TDM options to enable the smart SFP to encapsulate PDH (Plesiochronous Digital Hierarchy) and SDH (synchronous Digital Hierarchy) traffic over packet switched networks (PSNs).

## Options

|                                       |                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ces-psn-port-dmac-check-enable</b> | Enables checking of the destination MAC address of the incoming packets on the receiving SFP. If you have configured the destination MAC address using the <b>dmac-address</b> 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. |
| <b>sfp-type <i>sfp-type</i></b>       | Specify the smart SFP type you want to configure on the interface. Values: T1, E1, DS3, STM1, STM4, and STM16.                                                                                                                                                                                                                                                                    |
| <b>tdm-in-loop</b>                    | Enables looping back of the input path (Rx) of TDM traffic on the SFP TDM port.                                                                                                                                                                                                                                                                                                   |
| <b>tdm-out-loop</b>                   | Enables looping back of the output path (Tx) of TDM traffic on the SFP TDM port.                                                                                                                                                                                                                                                                                                  |

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 19.4.

## RELATED DOCUMENTATION

[show interfaces smart-sfp-statistics](#) | 1719

[show interfaces smart-sfp-defects](#) | 1705

## terminate (PPPoE Service Name Tables)

### IN THIS SECTION

- [Syntax | 1178](#)
- [Hierarchy Level | 1178](#)
- [Description | 1178](#)
- [Required Privilege Level | 1178](#)
- [Release Information | 1179](#)

### Syntax

```
terminate;
```

### Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name],
[edit protocols pppoe service-name-tables table-name service service-name agent-
specifier aci circuit-id-string ari remote-id-string]
```

### Description

Direct the router to immediately respond to a PPPoE Active Discovery Initiation (PADI) control packet received from a PPPoE client by sending the client a PPPoE Active Discovery Offer (PADO) packet. The PADO packet contains the name of the access concentrator (router) that can service the client request. The **terminate** action is the default action for a named service entry, **empty** service entry, **any** service entry, or agent circuit identifier/agent remote identifier (ACI/ARI) pair in a PPPoE service name table.

### 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 Release 10.0.

Support at `[edit protocols pppoe service-name-tables table-name service service-name agent-specifier aci circuit-id-string ari remote-id-string]` hierarchy level introduced in Junos OS Release 10.2.

## RELATED DOCUMENTATION

| [Configuring PPPoE Service Name Tables](#)

## thresholds

### IN THIS SECTION

- [Syntax | 1179](#)
- [Hierarchy Level | 1180](#)
- [Description | 1180](#)
- [Options | 1180](#)
- [Required Privilege Level | 1181](#)
- [Release Information | 1181](#)

## Syntax

```
thresholds {
 clear clear-value;
 interval interval-value;
 set set-value;
 warning-clear warning-clear-value;
 warning-set warning-set-value;
}
```

## Hierarchy Level

```
[edit interfaces interfaces-name link-degrade-monitor]
```

## Description

Configure the BER threshold values (such as set and clear thresholds) at which different corrective actions must be triggered on a degraded interface.

## Options

- |                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>clear</b> <i>clear-value</i>                 | The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted. You can configure this value in the $1E-n$ format, where 1 is the mantissa (remains constant) and $n$ is the exponent. For example, a threshold value of $1E-3$ refers to the BER threshold value of $1 \times 10^{-3}$ . The supported exponent range is 1 through 16, and the default value is 12.                                              |
| <b>interval</b> <i>interval-value</i>           | The number of consecutive link degrade events that are considered before any corrective action is taken. The supported value range for the interval is 1 through 256, and the default interval is 10.                                                                                                                                                                                                                                                                                        |
| <b>set</b> <i>set-value</i>                     | The BER threshold value at which the link is considered degraded and a corrective action, specified by the user, is triggered. You can configure this value in the $1E-n$ format, where 1 is the mantissa (remains constant) and $n$ is the exponent. For example, a threshold value of $1E-3$ refers to the BER threshold value of $1 \times 10^{-3}$ . The supported exponent range is 1 through 16, and the default value is 7.                                                           |
| <b>warning clear</b> <i>warning-clear-value</i> | The link clear warning threshold. Every time this threshold value is reached, a system message is logged to indicate that the link degrade condition has been cleared on the interface. You can configure this value in the $1E-n$ format, where 1 is the mantissa (remains constant) and $n$ is the exponent. For example, a threshold value of $1E-3$ refers to the BER threshold value of $1 \times 10^{-3}$ . The supported exponent range is 1 through 16, and the default value is 11. |
| <b>warning set</b> <i>warning-set-value</i>     | The link degrade warning threshold. Every time this threshold value is reached, a system message is logged to indicate that a link degrade has occurred on the interface. You can configure this value in the $1E-n$ format, where 1 is the mantissa (remains constant) and $n$ is the exponent. For example, a threshold value of $1E-3$ refers to the BER threshold value of $1 \times 10^{-3}$ . The supported exponent range is 1 through 16, and the default value is 9.                |

**NOTE:** The lower the BER with high confidence level, the longer it takes to estimate it. In such cases, a few packet drops might be noticed (based on the bit error distribution) before a link degrade event is detected.

## 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 Release 15.1.

## RELATED DOCUMENTATION

[Link Degrade Monitoring Overview | 600](#)

[link-degrade-monitor | 955](#)

[recovery | 1101](#)

[request interface link-degrade-recover | 1243](#)

## traceoptions (LLDP)

### IN THIS SECTION

- [Syntax | 1182](#)
- [Hierarchy Level | 1182](#)
- [Description | 1182](#)
- [Default | 1182](#)
- [Options | 1182](#)
- [Required Privilege Level | 1184](#)
- [Release Information | 1184](#)

## Syntax

```
traceoptions {
 file filename <files number> <size maximum-file-size> <world-readable | no-
world-readable>;
 flag flag <disable>;
}
```

## Hierarchy Level

```
[edit protocols lldp],
[edit routing-instances routing-instance-name protocols lldp]
```

## Description

Define tracing operations for the Link Layer Discovery Protocol (LLDP). You can trace messages under LLDP for LLDP and physical topology SNMP MIBs.

**NOTE:** The traceoptions statement is not supported on the QFX3000 QFabric system.

## Default

The default LLDP protocol-level trace options are inherited from the global **traceoptions** statement.

## Options

- disable** (Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.
- file *filename*** Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**. We recommend that you place spanning-tree protocol tracing output in the file **/var/log/stp-log**.
- files *number*** (Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

- **Range:** 2 through 1000 files
- **Default:** 1 trace file only

***flag***

Specify a tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **Values:** The following are the LLDP-specific tracing options:
  - **all**—Trace all operations.
  - **configuration**—Log configuration events.
  - **interface**—Trace interface update events.
  - **packet**—Trace packet events.
  - **protocol**—Trace protocol information.
  - **rtsock**—Trace socket events.
  - **snmp**—Trace SNMP configuration operations.
  - **vlan**—Trace VLAN update events.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events. This is the default. If you do not specify this option, only unusual or abnormal operations are traced.
- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.

- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

**no-world-readable** (Optional) Prevent any user from reading the log file. This is the default. If you do not include this option, tracing output is appended to an existing trace file.

**size maximum-file-size** (Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

- **Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB
- **Range:** 10 KB through the maximum file size supported on your system
- **Default:** 1 MB

**world-readable** (Optional) Allow any user to read the log file.

## Required Privilege Level

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced in Junos OS Release 9.0.

## RELATED DOCUMENTATION

[Configuring LLDP-MED \(CLI Procedure\)](#)

[Understanding LLDP and LLDP-MED on EX Series Switches](#)

[Understanding LLDP](#)

[Tracing LLDP Operations](#)

## traceoptions (Individual Interfaces)

### IN THIS SECTION

- [Syntax \(Individual interfaces with PTX Series, EX Series, ACX Series\) | 1185](#)
- [Syntax \(Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series\) | 1185](#)
- [Syntax \(OAMLFM with EX Series, QFX Series, NFX Series\) | 1186](#)
- [Syntax \(Interface process with ACX Series, SRX Series, MX Series, M Series, T Series\) | 1186](#)
- [Hierarchy Level \(Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series\) | 1186](#)
- [Hierarchy Level \(Interface process with ACX Series, SRX Series, MX Series, M Series, T Series\) | 1186](#)
- [Description | 1186](#)
- [Default | 1187](#)
- [Options | 1187](#)
- [Required Privilege Level | 1193](#)
- [Release Information | 1194](#)

### Syntax (Individual interfaces with PTX Series, EX Series, ACX Series)

```
traceoptions {
 file filename <files name> <size size> <world-readable | no-world-readable>;
 flag flag;
 match;
}
```

### Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series)

```
traceoptions {
 flag flag;
}
```

### Syntax (OAMLFM with EX Series, QFX Series, NFX Series)

```
traceoptions {
 file filename <files number> <match regex> <size size> <world-readable | no-
world-readable>;
 flag flag ;
 no-remote-trace;
}
```

### Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

```
traceoptions {
 file <filename> <files number> <match regular-expression> <size size> <world-
readable | no-world-readable>;
 flag flag <disable>;
 no-remote-trace;
}
```

### Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series)

```
[edit interfaces interface-name]
```

### Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series)

```
[edit interfaces]
```

## Description

Define tracing operations for individual interfaces.

To specify more than one tracing operation, include multiple **flag** statements.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** file in the directory */var/log/dcd*.

On EX Series, QFX Series, and NFX Series platforms, configure tracing options the link fault management.

On ACX Series, SRX Series, MX Series, M Series, and T Series platforms define tracing operations for the interface process (dcd).

## Default

If you do not include this statement, no interface-specific tracing operations are performed.

## Options

[Table 135 on page 1187](#) lists options for traceoption command for the following platforms:

**Table 135: Options for traceoptions**

| Option                      | Individual interfaces with PTX Series, ACX Series, EX Series                                                                                                                                                                                     | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series                                                                                                         | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series                                                                                                                                                                                 |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>file <i>filename</i></b> | —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log/dcd</code> . By default, interface process tracing output is placed in the file. |                                                                                    | —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log/dcd</code> . | —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log/dcd</code> . By default, interface process tracing output is placed in the file <b>dcd</b> . |

Table 135: Options for traceoptions (Continued)

| Option              | Individual interfaces with PTX Series, ACX Series, EX Series                                                                                                                                                                                                                                     | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series                                                                                                                                                                                                                                                                                                                                                                                                                           | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>files number</i> | <p>—(Optional)<br/>Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> |                                                                                    | <p>—(Optional)<br/>Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> | <p>—(Optional)<br/>Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><i>Range:</i> 2 through 1000</p> <p><i>Default:</i> 3 files</p> |

Table 135: Options for traceoptions (*Continued*)

| Option      | Individual interfaces with PTX Series, ACX Series, EX Series                                                                                                                                                                                                                                                                                                                                                                                            | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series                                                                                                                                                                                                                                                                                                                                                                      | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series                                                                                                                                                                                                                                                                                                                                                                                                                                             | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>flag</i> | <p>–Tracing operation to perform. To specify more than one tracing operation, include multiple <i>flag</i> statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"> <li>• <b>all</b>—All interface tracing operations</li> <li>• <b>event</b>—Interface events</li> <li>• <b>ipc</b>—Interface interprocess communication (IPC) messages</li> <li>• <b>media</b>—Interface media changes</li> </ul> | <p>–Tracing operation to perform. To specify more than one tracing operation, include multiple <i>flag</i> statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"> <li>• <b>all</b>—All interface tracing operations</li> <li>• <b>event</b>—Interface events</li> <li>• <b>ipc</b>—Interface interprocess communication (IPC) messages</li> <li>• <b>media</b>—Interface media changes</li> </ul> | <p>–Tracing operation to perform. To specify more than one tracing operation, include multiple <i>flag</i> statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li>• <b>action-profile</b>—Trace action profile invocation events.</li> <li>• <b>all</b>—Trace all events.</li> <li>• <b>configuration</b>—Trace configuration events.</li> <li>• <b>protocol</b>—Trace protocol processing events.</li> <li>• <b>routing socket</b>—Trace routing socket events.</li> </ul> | <p>–Tracing operation to perform. To specify more than one tracing operation, include multiple <i>flag</i> statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li>• <b>all</b></li> <li>• <b>change-events</b>—Log changes that produce configuration events</li> <li>• <b>config-states</b>—Log the configuration state machine changes</li> <li>• <b>kernel</b>—Log configuration IPC messages to kernel</li> <li>• <b>kernel-detail</b>—Log details of</li> </ul> |

Table 135: Options for traceoptions (*Continued*)

| Option       | Individual interfaces with PTX Series, ACX Series, EX Series                                                                           | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series                                                     | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series                           | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
|              | <ul style="list-style-type: none"> <li>• <b>q921</b>—Trace ISDN Q.921 frames</li> <li>• <b>q931</b>—Trace ISDN Q.931 frames</li> </ul> | <ul style="list-style-type: none"> <li>• <b>q921</b>—Trace ISDN Q.921 frames</li> <li>• <b>q931</b>—Trace ISDN Q.931 frames</li> </ul> |                                                                                                | configuration messages to kernel                                             |
| <b>match</b> | —(Optional) Regular expression for lines to be traced.                                                                                 |                                                                                                                                        | —(Optional) Refine the output to log only those lines that match the given regular expression. |                                                                              |

Table 135: Options for traceoptions (Continued)

| Option           | Individual interfaces with PTX Series, ACX Series, EX Series                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| size <i>size</i> | <p>—(Optional)<br/>Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches this size, it is renamed <b>trace-file.0</b>. When the <b>trace-file</b> again reaches its maximum size, <b>trace-file.0</b> is renamed <b>trace-file.1</b> and <b>trace-file</b> is renamed <b>trace-file.0</b>. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> |                                                                                    | <p>—(Optional)<br/>Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option.</p> <p><i>Syntax:</i> <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB</p> | <p>—(Optional)<br/>Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches this size, it is renamed <b>trace-file.0</b>. When the <b>trace-file</b> again reaches its maximum size, <b>trace-file.0</b> is renamed <b>trace-file.1</b> and <b>trace-file</b> is renamed <b>trace-file.0</b>. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number</p> |

Table 135: Options for traceoptions (*Continued*)

| Option                   | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series                                                                                                                  | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series                                                                                                                                                                                                              |
|--------------------------|--------------------------------------------------------------|------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                          |                                                              |                                                                                    | <p><i>Range:</i> 10 KB through 1 GB</p> <p><i>Default:</i> 128 KB</p> <p><i>Default:</i> If you do not include this option, tracing output is appended to an existing trace file.</p> | <p>of trace files with the <b>files</b> option.</p> <p><i>Syntax:</i> <b>xk</b> to specify kilobytes, <b>xm</b> to specify megabytes, or <b>xg</b> to specify gigabytes</p> <p><i>Range:</i> 10 KB through the maximum file size supported on your router</p> <p><i>Default:</i> 1 MB</p> |
| <b>no-world-readable</b> | —(Optional) Prevent any user from reading the log file.      |                                                                                    | —(Optional) Restrict file access to the user who created the file.                                                                                                                    | —(Optional) Disallow any user to read the log file.                                                                                                                                                                                                                                       |
| <b>world-readable</b>    | —(Optional) Allow any user to read the log file.             |                                                                                    | —(Optional) Enable unrestricted file access.                                                                                                                                          | —(Optional) Allow any user to read the log file.                                                                                                                                                                                                                                          |

Table 135: Options for traceoptions (*Continued*)

| Option                    | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series                                                                                                        |
|---------------------------|--------------------------------------------------------------|------------------------------------------------------------------------------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>disable</b>            |                                                              |                                                                                    |                                                                      | –(Optional)<br>Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b> . |
| <b>no-remote-trace</b>    |                                                              |                                                                                    | –(Optional)<br>Disable the remote trace.                             | -                                                                                                                                                                                   |
| <b>match <i>regex</i></b> |                                                              |                                                                                    |                                                                      | –(Optional) Refine the output to include only those lines that match the given regular expression.                                                                                  |

### Required Privilege Level

interface—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

## RELATED DOCUMENTATION

[Tracing Operations of an Individual Router Interface](#)

[Tracing Operations of an Individual Router or Switch Interface](#)

*Example: Configuring Ethernet OAM Link Fault Management*

*Configuring Ethernet OAM Link Fault Management*

*Tracing Operations of the Interface Process*

## traceoptions (LACP)

### IN THIS SECTION

- [Syntax | 1194](#)
- [Hierarchy Level | 1195](#)
- [Description | 1195](#)
- [Default | 1195](#)
- [Options | 1195](#)
- [Required Privilege Level | 1196](#)
- [Release Information | 1196](#)

## Syntax

```
traceoptions {
 file <filename> <files number> <size size> <world-readable | no-world-
readable>;
 flag flag;
```

```
no-remote-trace;
}
```

## Hierarchy Level

```
[edit protocols lacp]
```

## Description

Define tracing operations for the LACP protocol.

## Default

If you do not include this statement, no LACP protocol tracing operations are performed.

## Options

**filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, interface process tracing output is placed in the file `lacpd`.

**files number**—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000
- **Default:** 3 files

**flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events

- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option:

- **Syntax:** ***xk*** to specify kilobytes, ***xm*** to specify megabytes, or ***xg*** to specify gigabytes
- **Range:** 10 KB through the maximum file size supported on your router
- **Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

## 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 Release 7.6.

## RELATED DOCUMENTATION

| [Configuring LACP for Aggregated Ethernet Interfaces](#)

## traceoptions (PPPoE)

### IN THIS SECTION

- [Syntax | 1197](#)
- [Hierarchy Level | 1197](#)
- [Description | 1197](#)
- [Options | 1198](#)
- [Required Privilege Level | 1200](#)
- [Release Information | 1200](#)

### Syntax

```
traceoptions {
 file <filename> <files number> <match regular-expression > <size maximum-
file-size> <world-readable | no-world-readable>;
 filter {
 aci regular-expression;
 ari regular-expression;
 service-name regular-expression;
 underlying-interface interface-name;
 }
 flag flag;
 level (all | error | info | notice | verbose | warning);
 no-remote-trace;
}
```

### Hierarchy Level

```
[edit protocols pppoe]
```

### Description

Define tracing operations for PPPoE processes.

## Options

**file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

**files *number***—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

- **Range:** 2 through 1000
- **Default:** 3 files

**disable**—Disable this trace flag.

**filter**—Additional filter to refine the output to display particular subscribers. Filtering based on the following subscriber identifiers simplifies troubleshooting in a scaled environment.

**BEST PRACTICE:** Due to the complexity of agent circuit identifiers and agent remote identifiers, we recommend that you do not try an exact match when filtering on these options. For service names, searching on the exact name is appropriate, but you can also use a regular expression with that option.

- **aci *regular-expression***—Regular expression to match the agent circuit identifier provided by PPPoE client.
- **ari *regular-expression***—Regular expression to match the agent remote identifier provided by PPPoE client.
- **service *regular-expression***—Regular expression to match the name of PPPoE service.
- **underlying-interface *interface-name***—Name of a PPPoE underlying interface. You cannot use a regular expression for this filter option.

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.
- **config**—Trace configuration events.
- **events**—Trace events.
- **gres**—Trace GRES events.
- **init**—Trace initialization events.

- **interface-db**—Trace interface database operations.
- **memory**—Trace memory processing events.
- **protocol**—Trace protocol events.
- **rtsock**—Trace routing socket events.
- **session-db**—Trace connection events and flow.
- **signal**—Trace signal operations.
- **state**—Trace state handling events.
- **timer**—Trace timer processing.
- **ui**—Trace user interface processing.

**level**—Level of tracing to perform. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match notice messages about conditions requiring special handling.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.
- **Default: error**

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**no-remote-trace**—Disable remote tracing.

**no-world-readable**—(Optional) Disable unrestricted file access.

**size *maximum-file-size***—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

- **Syntax:** *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB
- **Range:** 10240 through 1073741824
- **Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

### Required Privilege Level

**trace**—To view this statement in the configuration.

**trace-control**—To add this statement to the configuration.

### Release Information

Statement introduced in Junos OS Release 9.6.

Option **filter** introduced in Junos OS Release 12.3

### RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

[Tracing PPPoE Operations | 57](#)

## tx-duration

### IN THIS SECTION

- [Syntax | 1201](#)
- [Hierarchy Level | 1201](#)
- [Description | 1201](#)
- [Options | 1201](#)
- [Required Privilege Level | 1201](#)
- [Release Information | 1201](#)

## Syntax

```
expected-defect {
 tx-duration ;
}
```

## Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

## Description

The expected duration for which the peer MEP should suppress the LoC alarms.

## Options

**Minimum value** The minimum value at which the peer MEP should suppress the LoC alarms is 120 *seconds*.

**Minimum value** The maximum acceptable value at which the peer MEP should suppress the LoC alarms is 3600 *seconds*.

**Default** 900 *seconds*.

## 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 Release 19.1.

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | 830

---

*show oam ethernet connectivity-fault-management mep-database*

## tx-enable

### IN THIS SECTION

- [Syntax | 1202](#)
- [Hierarchy Level | 1202](#)
- [Description | 1202](#)
- [Default | 1202](#)
- [Required Privilege Level | 1202](#)
- [Release Information | 1203](#)

### Syntax

```
expected-defect {
 tx-enable ;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

### Description

Enable the ethernet expected defect (ETH-ED) function to control if EDM transmission need to be triggered on ISSU.

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

### Default

The MEP does not generate EDM PDUs by default.

### 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 Release 19.1.

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | 830

*show oam ethernet connectivity-fault-management mep-database*

## underlying-interface

### IN THIS SECTION

- [Syntax](#) | 1203
- [Hierarchy Level](#) | 1203
- [Description](#) | 1204
- [Options](#) | 1204
- [Required Privilege Level](#) | 1204
- [Release Information](#) | 1204

## Syntax

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

## Hierarchy Level

```
[edit interfaces pp0 unit logical-unit-number pppoe-options],
[edit interfaces demux0 unit logical-unit-number demux-options],
[edit logical-systems logical-system-name interfaces demux0 unit logical-unit-
number demux-options],
```

```
[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces demux0 unit logical-unit-number demux-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name interfaces pp0 unit logical-unit-number pppoe-options]
```

## Description

Configure the interface on which PPP over Ethernet is running.

For demux interfaces, configure the underlying interface on which the demultiplexing (demux) interface is running.

## Options

***interface-name***—Name of the interface on which PPP over Ethernet or demux is running. For example, **at-0/0/1.0** (ATM VC), **fe-1/0/1.0** (Fast Ethernet interface), **ge-2/0/0.0** (Gigabit Ethernet interface), **ae1.0** (for IP demux on an aggregated Ethernet interface), or **ae1** (for VLAN demux on an aggregated Ethernet interface).

**NOTE:** Demux interfaces are currently supported on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet interfaces, or aggregated Ethernet devices.

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Support for aggregated Ethernet added in Junos OS Release 9.4.

## RELATED DOCUMENTATION

[Configuring an IP Demultiplexing Interface](#)

Configuring a VLAN Demultiplexing Interface

Configuring the PPPoE Underlying Interface

## unit

### IN THIS SECTION

- Syntax | 1205
- Hierarchy Level | 1214
- Description | 1214
- Options | 1215
- Required Privilege Level | 1215
- Release Information | 1215

## Syntax

```
unit logical-unit-number {
 accept-source-mac {
 mac-address mac-address {
 policer {
 input cos-policer-name;
 output cos-policer-name;
 }
 }
 }
 accounting-profile name;
 advisory-options {
 downstream-rate rate;
 upstream-rate rate;
 }
 allow-any-vci;
 atm-scheduler-map (map-name | default);
 auto-configure {
 agent-circuit-identifier {
 dynamic-profile profile-name;
```

```

}
line-identity {
 include {
 accept-no-ids;
 circuit-id;
 remote-id;

 }
 dynamic-profile profile-name;
}
}
backup-options {
 interface interface-name;
}
bandwidth rate;
cell-bundle-size cells;
clear-dont-fragment-bit;
compression {
 rtp {
 maximum-contexts number <force>;
 f-max-period number;
 queues [queue-numbers];
 port {
 minimum port-number;
 maximum port-number;
 }
 }
}
compression-device interface-name;
copy-tos-to-outer-ip-header;
demux {
 inet {
 address-source address;
 auto-configure {
 address-ranges {
 authentication {
 password password-string;
 username-include {
 auth-server-realm realm-string;
 delimiter delimiter-character;
 domain-name domain-name;
 interface-name;
 source-address;
 }
 }
 }
 }
 }
}

```



```

demux-destination family;
demux-source family;
demux-options {
 underlying-interface interface-name;
}
description text;
etree-ac-role (leaf | root);
interface {
 l2tp-interface-id name;
 (dedicated | shared);
}
dialer-options {
 activation-delay seconds;
 callback;
 callback-wait-period time;
 deactivation-delay seconds;
 dial-string [dial-string-numbers];
 idle-timeout seconds;
 incoming-map {
 caller caller-id | accept-all;
 initial-route-check seconds;
 load-interval seconds;
 load-threshold percent;
 pool pool-name;
 redial-delay time;
 watch-list {
 [routes];
 }
 }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
 activation-priority priority;
 bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
family family-name {
 ... the family subhierarchy appears after the main [edit interfaces
interface-name unit logical-unit-number] hierarchy ...

```

```

}
fragment-threshold bytes;
host-prefix-only;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
 (pop | pop-pop | pop-swap | push | push-push | swap |
 swap-push | swap-swap);
 inner-tag-protocol-id tpid;
 inner-vlan-id number;
 tag-protocol-id tpid;
 vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
 input-policer policer-name;
 input-three-color policer-name;
 output-policer policer-name;
 output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
 up-count cells;
 down-count cells;
}
oam-period (disable | seconds);
output-vlan-map {
 (pop | pop-pop | pop-swap | push | push-push | swap |
 swap-push | swap-swap);
 inner-tag-protocol-id tpid;
 inner-vlan-id number;
 tag-protocol-id tpid;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;

```

```

ppp-options {
 mru size;
 mtu (size | use-lower-layer);
 chap {
 access-profile name;
 default-chap-secret name;
 local-name name;
 passive;
 }
 compression {
 acfc;
 pfc;
 }
 dynamic-profile profile-name;
 ipcp-suggest-dns-option;
 lcp-restart-timer milliseconds;
 loopback-clear-timer seconds;
 ncp-restart-timer milliseconds;
 pap {
 access-profile name;
 default-pap-password password;
 local-name name;
 local-password password;
 passive;
 }
}
pppoe-options {
 access-concentrator name;
 auto-reconnect seconds;
 (client | server);
 service-name name;
 underlying-interface interface-name;
}
pppoe-underlying-options {
 access-concentrator name;
 direct-connect;
 dynamic-profile profile-name;
 max-sessions number;
}
proxy-arp;
service-domain (inside | outside);
shaping {
 (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate

```

```

sustained rate burst length);
 queue-length number;
}
short-sequence;
targeted-distribution;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
 backup-destination address;
 destination address;
 key number;
 routing-instance {
 destination routing-instance-name;
 }
 source source-address;
 ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
family family {
 accounting {
 destination-class-usage;
 source-class-usage {
 (input | output | input output);
 }
 }
 access-concentrator name;
 address address {
 ... the address subhierarchy appears after the main [edit interfaces
interface-name unit logical-unit-number family family-name] hierarchy ...
 }
 bundle interface-name;
 core-facing;
 demux-destination {
 destination-prefix;
 }
 demux-source {

```

```
 source-prefix;
}
direct-connect;
duplicate-protection;
dynamic-profile profile-name;
filter {
 group filter-group-number;
 input filter-name;
 input-list [filter-names];
 output filter-name;
 output-list [filter-names];
}
interface-mode (access | trunk);
ipsec-sa sa-name;
keep-address-and-control;
mac-validate (loose | strict);
max-sessions number;
mtu bytes;
multicast-only;
no-redirects;
policer {
 arp policer-template-name;
 input policer-template-name;
 output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
 fail-filter filter-name
 mode loose;
}
sampling {
 input;
 output;
}
service {
 input {
 post-service-filter filter-name;
 service-set service-set-name <service-filter filter-name>;
 }
}
```

```

 }
 output {
 service-set service-set-name <service-filter filter-name>;
 }
}
service-name-table table-name
targeted-options {
 backup backup;
 group group;
 primary primary;
 weight ($junos-interface-target-weight | weight-value);
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address destination-
profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
address address {
 arp ip-address (mac | multicast-mac) mac-address <publish>;
 broadcast address;
 destination address;
 destination-profile name;
 eui-64;
 primary-only;
 multipoint-destination address {
 dlci dlci-identifier;
 epd-threshold cells <plp1 cells>;
 inverse-arp;
 oam-liveness {
 up-count cells;
 down-count cells;
 }
 oam-period (disable | seconds);
 shaping {
 (cbr rate | rtvbr burst length peak rate sustained rate |
vbr burst length peak rate sustained rate);
 queue-length number;
 }
 vci vpi-identifier.vci-identifier;
 }
 preferred;
}

```



## Options

*logical-unit-number*—Number of the logical unit.

- **Range:** 0 through 1,073,741,823 for demux, PPPoE, and pseudowire static interfaces. 0 through 16,385 for all other static interface types.

**etree-ac-role (leaf | root)**—To configure an interface as either leaf or root.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

Range increased for static pseudowire interfaces to 1,073,741,823 in Junos OS Release 18.3R1.

## RELATED DOCUMENTATION

[Configuring Logical Interface Properties](#)

[Junos OS Services Interfaces Library for Routing Devices](#)

## unnumbered-address (Dynamic Profiles)

### IN THIS SECTION

- [Syntax](#) | 1216
- [Hierarchy Level](#) | 1216
- [Description](#) | 1216
- [Options](#) | 1216

- Required Privilege Level | 1218
- Release Information | 1218

## Syntax

```
unnumbered-address interface-name <preferred-source-address address>;
```

## Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family],
[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family]
```

## Description

For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the **\$junos-loopback-interface-address** predefined variable.

You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.

## Options

***interface-name*** Name of the interface from which the local address is derived. The specified interface must have a logical unit number, a configured IP address, and must not be an unnumbered interface. This value can be a specific interface name or the **\$junos-loopback-interface** predefined variable.

When defining the **unnumbered-address** statement using a static interface, keep the following in mind:

- If you choose to include the **routing-instance** statement at the [edit dynamic-profiles] hierarchy level, that statement must be configured with a dynamic value by using the **\$junos-routing-instance** predefined variable. In addition, whatever static unnumbered interface you specify must belong to that routing instance; otherwise, the profile instantiation fails.
- If you choose to not include the **routing-instance** statement at the [edit dynamic-profiles] hierarchy level, the unnumbered-address statement uses the default routing instance. The use of the default routing instance requires that the unnumbered interface be configured statically and that it reside in the default routing instance.

**NOTE:** When you specify a static logical interface for the unnumbered interface in a dynamic profile that includes the **\$junos-routing-instance** predefined variable, you must not configure a preferred source address, whether with the **\$junos-preferred-source-address** predefined variable, the **\$junos-preferred-source-ipv6-address** predefined variable, or the **preferred-source-address** statement. Configuring the preferred source address in this circumstance causes a commit failure.

When defining the **unnumbered-address** statement using the **\$junos-loopback-interface** predefined variable, keep the following in mind:

- To use the **\$junos-loopback-interface** predefined variable, the dynamic profile must also contain the **routing-instance** statement configured with the **\$junos-routing-instance** predefined variable at the [edit dynamic-profiles] hierarchy level.
- The applied loopback interface is based on the dynamically obtained routing instance of the subscriber.

*address*

(Optional) Secondary IP address of the donor interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. This value can be a static IP address, the **\$junos-preferred-source-address** predefined variable for the inet family, or the **\$junos-preferred-source-ipv6-address** predefined variable for the inet6 family.

When defining the **preferred-source-address** value using a static IP address, keep the following in mind:

- The unnumbered interface must be statically configured.
- The IP address specified as the **preferred-source-address** must be configured in the specified unnumbered interface.

When defining the **preferred-source-address** value using the **\$junos-preferred-source-address** or the **\$junos-preferred-source-ipv6-address** predefined variables, keep the following in mind:

- You must configure the **unnumbered-address** statement using the **\$junos-loopback-interface** predefined variable.
- You must configure the **routing-instance** statement using the **\$junos-routing-instance** predefined variable at the **[edit dynamic-profiles]** hierarchy level.
- The preferred source address chosen is based on the dynamically applied loopback address which is in turn derived from the dynamically obtained routing instance of the subscriber. The configured loopback address with the closest network match to the user IP address is selected as the preferred source address.

## 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 Release 9.2.

Support for the **\$junos-preferred-source-address** and **\$junos-preferred-source-ipv6-address** predefined variables introduced in Junos OS Release 9.6.

Support for the **\$junos-loopback-interface** predefined variable introduced in Junos OS Release 9.6.

## RELATED DOCUMENTATION

| [Dynamic Profiles Overview](#)

## unnumbered-address (PPP)

### IN THIS SECTION

- [Syntax | 1219](#)
- [Hierarchy Level | 1219](#)
- [Description | 1219](#)
- [Options | 1219](#)
- [Required Privilege Level | 1220](#)
- [Release Information | 1220](#)

### Syntax

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

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit logical-systems logical-system-name interfaces interface-name unit logical-
unit-number family inet]
```

### Description

For interfaces with PPP encapsulation, enable the local address to be derived from the specified interface.

### Options

***interface-name***—Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address.

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

## Required Privilege Level

interface—To view this statement in the configuration.

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

## Release Information

Statement introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

[Configuring IPCP Options for Interfaces with PPP Encapsulation](#)

## Unused

### IN THIS SECTION

- [Syntax | 1220](#)
- [Hierarchy Level | 1221](#)
- [Description | 1221](#)
- [Options | 1221](#)
- [Required Privilege Level | 1221](#)
- [Release Information | 1221](#)

## 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 be 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 than 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.

Refer to "[Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers](#)" on page 418 for more information.

(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-name* 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

[Configuring Port Speed on Routing Devices and Line Cards | 418](#)

[speed \(Ethernet\) | 1137](#)

[number-of-sub-ports | 1042](#)

[PTX10003 Port Panel](#)

[Configuring the Port Speed on the JNP10K-LC1201 by Using New Port Profile Configuration | 452](#)

## virtual-control-channel

### IN THIS SECTION

- [Syntax | 1222](#)
- [Hierarchy Level | 1223](#)
- [Description | 1223](#)
- [Options | 1223](#)
- [Required Privilege Level | 1223](#)
- [Release Information | 1223](#)

## Syntax

```
virtual-control-channel channel-name {
 west-interface name;
 east-interface name;
}
```

## Hierarchy Level

```
[edit protocols protection-group ethernet-ring name (east-interface | west-interface)]
```

## Description

Specify virtual control channels which are logical interfaces on the east and west interfaces of the major ring.

## Options

**west-interface** *name*—Logical interface on the major ring's west port.

**east-interface** *name*—Logical interface on the major ring's east port.

## 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 Release 14.2.

## RELATED DOCUMENTATION

[Ethernet Ring Protection Switching Overview](#)

[Configuring Ethernet Ring Protection Switching on Switches \(CLI Procedure\)](#)

## virtual-switch

### IN THIS SECTION

- [Syntax | 1224](#)
- [Hierarchy Level | 1224](#)
- [Description | 1224](#)
- [Required Privilege Level | 1224](#)
- [Release Information | 1224](#)

### Syntax

```
virtual-switch name bridge-domain name vlan-id [vlan-ids];
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name default-x]
```

### Description

Specify the routing-instance type as a virtual switch, under which bridge-domain MIPs must be enabled.

### 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 Release 9.6.

## RELATED DOCUMENTATION

| *Configuring Maintenance Intermediate Points (MIPs)*

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

## IN THIS SECTION

- [Syntax | 1225](#)
- [Hierarchy Level | 1225](#)
- [Description | 1225](#)
- [Options | 1226](#)
- [Required Privilege Level | 1226](#)
- [Release Information | 1226](#)

## Syntax

```
vlan-rule (high-low | odd-even);
```

## Hierarchy Level

```
[edit chassis fpc slot pic slot forwarding-mode vlan-steering]
```

## Description

Configure the interoperation mode of the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) when interoperating with 100 gigabit Ethernet interfaces from other vendors.

If no VLAN rule is configured, all tagged packets are distributed to PFE0.

## Options

- high-low** VLAN IDs 1 through 2047 are distributed to PFE0 and VLAN IDs 2048 through 4096 are distributed to PFE1.
- odd-even** Odd number VLAN IDs are distributed to PFE1 and even number VLAN IDs are distributed to PFE0.

## 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 Release 10.4.

## RELATED DOCUMENTATION

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 231](#)

[forwarding-mode \(100-Gigabit Ethernet\) | 886](#)

[vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1226](#)

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

### IN THIS SECTION

- [Syntax | 1227](#)
- [Hierarchy Level | 1227](#)
- [Description | 1227](#)
- [Required Privilege Level | 1227](#)
- [Release Information | 1227](#)

## Syntax

```
vlan-steering {
 vlan-rule (high-low | odd-even);
}
```

## Hierarchy Level

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

## Description

Configure the 100-Gigabit Ethernet Type 4 PIC with CFP (PD-1CE-CFP-FPC4) to interoperate with 100 gigabit Ethernet interfaces from other vendors.

The other statement is explained separately.

## 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 Release 9.4.

## RELATED DOCUMENTATION

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 231](#)

[forwarding-mode \(100-Gigabit Ethernet\) | 886](#)

[sa-multicast \(100-Gigabit Ethernet\) | 1113](#)

[vlan-rule \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1225](#)

# Operational Commands

## IN THIS CHAPTER

- clear interfaces interface-set statistics | **1229**
- clear interfaces interval | **1231**
- clear interfaces aeX forwarding-options load-balance state | **1233**
- clear interfaces aggregate forwarding-options load-balance state | **1235**
- clear interfaces transport pm | **1236**
- clear protection-group ethernet-ring statistics | **1238**
- prbs-test-start | **1239**
- prbs-test-stop | **1242**
- request interface link-degrade-recover | **1243**
- request interface mc-ae switchover (Multichassis Link Aggregation) | **1247**
- request interface (revert | switchover) (Aggregated Ethernet Link Protection) | **1250**
- request lacp link-switchover | **1251**
- show chassis hardware | **1253**
- show chassis pic | **1274**
- show ethernet-switching redundancy-groups | **1311**
- show interfaces (Adaptive Services) | **1317**
- show interfaces (Aggregated Ethernet) | **1328**
- show interfaces demux0 (Demux Interfaces) | **1352**
- show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) | **1368**
- show interfaces (far-end-interval) | **1432**
- show interfaces (Fast Ethernet) | **1434**
- show interfaces | **1463**
- show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) | **1602**
- show interfaces (PPPoE) | **1627**

- [show interfaces interface-set \(Ethernet Interface Set\) | 1645](#)
- [show interfaces interface-set queue | 1651](#)
- [show interfaces interval | 1663](#)
- [show interfaces irb | 1667](#)
- [show interfaces mac-database | 1680](#)
- [show interfaces mc-ae | 1688](#)
- [show interfaces port-profile | 1695](#)
- [show interfaces prbs-stats | 1703](#)
- [show interfaces smart-sfp-defects | 1705](#)
- [show interfaces smart-sfp-statistics | 1719](#)
- [show interfaces transport pm | 1737](#)
- [show l2-learning instance | 1747](#)
- [show l2-learning redundancy-groups | 1750](#)
- [show lacp interfaces | 1756](#)
- [show pppoe interfaces | 1765](#)
- [show pppoe service-name-tables | 1771](#)
- [show pppoe sessions | 1776](#)
- [show pppoe statistics | 1778](#)
- [show pppoe underlying-interfaces | 1782](#)
- [show pppoe version | 1793](#)
- [traceroute ethernet | 1795](#)

## clear interfaces interface-set statistics

### IN THIS SECTION

- [Syntax | 1230](#)
- [Description | 1230](#)
- [Options | 1230](#)
- [Required Privilege Level | 1230](#)

- [Output Fields | 1230](#)
- [Sample Output | 1230](#)
- [Release Information | 1231](#)

## Syntax

```
clear interfaces interface-set statistics interface-set-name
```

## Description

Set interface set statistics to zero.

## Options

***interface-set-name*** Set statistics on a specified interface set to zero. Wildcard values can be used in the interface set name. This command will not clear the statistics of the member logical interfaces.

## Required Privilege Level

clear

## Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**clear interfaces interface-set statistics**

```
user@host> clear interfaces interface-set statistics
```

## Release Information

Command introduced in Junos OS Release 8.5.

## clear interfaces interval

### IN THIS SECTION

- [Syntax | 1231](#)
- [Description | 1231](#)
- [Options | 1231](#)
- [Required Privilege Level | 1231](#)
- [Output Fields | 1232](#)
- [Sample Output | 1232](#)
- [Release Information | 1233](#)

## Syntax

```
clear interfaces interval interface-name
```

## Description

Clear the channel service unit (CSU) alarm and defect counters so that only the current time interval is displayed. This operation affects the **show interface interval** command, but not an SNMP query.

## Options

*interface-name*                      Name of a particular interface.

## Required Privilege Level

clear

## Output Fields

See "[show interfaces interval](#)" on page 1663 for an explanation of output fields.

## Sample Output

### clear interfaces interval

The following example displays the output for a T3 interface before and after the **clear interfaces** command is entered:

```
user@host> show interfaces interval t3-0/3/0:4
Physical interface: t3-0/3/0:4, SNMP ifIndex: 23
 17:43-current:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
 17:28-17:43:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
 17:13-17:28:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
 16:58-17:13:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
 16:43-16:58:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
 16:28-16:43:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
CES: 195, CSES: 195, SEFS: 195, UAS: 206
 14:58-15:13:
 LCV: 35, PCV: 163394, CCV: 54485, LES: 0, PES: 35, PSES: 35, CES:
35, CSES: 35, SEFS: 35, UAS: 32
Interval Total:
 LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
CES: 230, CSES: 230, SEFS: 230, UAS: 238
```

```
user@host> clear interfaces interval t3-0/3/0:4
```

```
user@host> show interfaces interval t3-0/3/0:4
Physical interface: t3-0/3/0:4, SNMP ifIndex: 23
 17:43-current:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
 Interval Total:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0, SEFS: 0,
UAS: 0
```

## Release Information

Command introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [show interfaces interval](#) | [1663](#)

## clear interfaces aeX forwarding-options load-balance state

### IN THIS SECTION

- [Syntax](#) | [1234](#)
- [Description](#) | [1234](#)
- [Options](#) | [1234](#)
- [Required Privilege Level](#) | [1234](#)
- [Sample Output](#) | [1234](#)
- [Release Information](#) | [1234](#)

## Syntax

```
clear interfaces aeX unit logical-unit-number aggregate forwarding-options load-
balance state
```

## Description

Clear the specified aggregate Ethernet interface load balancing state and re-create it newly. 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.

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.

## Options

|                                              |                                                                          |
|----------------------------------------------|--------------------------------------------------------------------------|
| <b>aeX</b>                                   | Name of a particular aggregated Ethernet interface.                      |
| <b><i>logical-unit-number</i></b>            | Number of the logical unit of the interface.                             |
| <b>forwarding-options load-balance state</b> | Cause the load-balancing state to be cleared for the specific interface. |

## Required Privilege Level

clear

## Sample Output

```
clear interfaces aeX aggregate forwarding-options
```

```
user@host> clear interfaces ae1 aggregate forwarding-options load-balance state
```

## Release Information

Command introduced in Junos OS Release 13.2R1.

## RELATED DOCUMENTATION

[show interfaces interval](#) | 1663

# clear interfaces aggregate forwarding-options load-balance state

## IN THIS SECTION

- [Syntax](#) | 1235
- [Description](#) | 1235
- [Options](#) | 1235
- [Required Privilege Level](#) | 1235
- [Sample Output](#) | 1236
- [Release Information](#) | 1236

## Syntax

```
clear interfaces aggregate forwarding-options load-balance state
```

## Description

Clear all the aggregate Ethernet interface load balancing states and re-create them newly. 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.

## Options

*interface-name*                      Name of a particular interface.

## Required Privilege Level

clear

## Sample Output

### clear interfaces aggregate forwarding-options

```
user@host> clear interfaces aggregate forwarding-options load-balance state
```

## Release Information

Command introduced in Junos OS Release 13.2R1.

## RELATED DOCUMENTATION

| [show interfaces interval](#) | [1663](#)

## clear interfaces transport pm

### IN THIS SECTION

- [Syntax](#) | [1236](#)
- [Description](#) | [1237](#)
- [Options](#) | [1237](#)
- [Required Privilege Level](#) | [1237](#)
- [Output Fields](#) | [1237](#)
- [Sample Output](#) | [1237](#)
- [Release Information](#) | [1237](#)

## Syntax

```
clear interfaces transport pm (all | optics | otn) (all | current | current-day)
(all | interface-name)
```

## Description

Clear optics and OTN information from the transport performance monitoring data.

## Options

(all | optics | otn)—Clear both optics and OTN information or either only optics or only OTN information.

(all | current | current-day)—Clear information for the current 15-minute interval, the ninety-six 15-minute intervals, the current day, and the previous day; information only for the current 15-minute interval; or information only for the current 24 hours.

(all | *interface-name*)—Clear information for all interfaces or only for the specified interface (for example, *et-fpc/ pic/ port*).

## Required Privilege Level

clear

## Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**clear interfaces transport pm**

```
user@host> clear interfaces transport pm transport otn current all
```

## Release Information

Command introduced in Junos OS Release 14.2.

## RELATED DOCUMENTATION

---

[show interfaces transport pm | 1737](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 459](#)

---

[tca | 757](#)

## clear protection-group ethernet-ring statistics

### IN THIS SECTION

- [Syntax | 1238](#)
- [Description | 1238](#)
- [Options | 1238](#)
- [Required Privilege Level | 1238](#)
- [Output Fields | 1239](#)
- [Sample Output | 1239](#)
- [Sample Output | 1239](#)
- [Release Information | 1239](#)

### Syntax

```
clear protection-group ethernet-ring statistics
<group-name group-name>
```

### Description

On MX Series routers, clear the statistics for all Ethernet ring protection groups or a specific Ethernet ring protection group.

### Options

**group-name** *group-name* (Optional) Clear the Ethernet ring protection statistics for the specified group.

### Required Privilege Level

view

## Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear protection-group ethernet-ring statistics

To clear all Ethernet ring protection group statistics for all protection groups, use the following command:

```
user@host> clear protection-group ethernet-ring statistics
```

## Sample Output

### clear protection-group ethernet-ring statistics

To clear Ethernet ring protection group statistics for the group *my\_prot\_group*, use the following command:

```
user@host> clear protection-group ethernet-ring statistics group-name my_prot_group
```

## Release Information

Command introduced in Junos OS Release 9.4.

## prbs-test-start

### IN THIS SECTION

- [Syntax | 1240](#)
- [Description | 1240](#)
- [Options | 1240](#)

- Required Privilege Level | 1241
- Release Information | 1241

## Syntax

```
test interfaces ifd-name prbs-test-start pattern-type type direction 0/1 flip
0/1
```

## Description

You can check the physical link connectivity by issuing this command that starts collecting the Pseudo Random Binary Sequence (PRBS) statistics. The PRBS pattern generation and verification validates the physical link connectivity in a routers. If the PRBS test passes with PRBS31 pattern type consistently, it indicates that the quality of signal received is good.

This command provides the PRBS test statistics while test is in progress. Use the **clear interfaces statistics** command to clear the collected statistics.

Use the **show interfaces interface-name prbs-stats** command to view the collected statistics.

### NOTE:

- While running PRBS statistics, the link will be down.
- The interface link status goes down when PRBS is enabled and the interface state is back to the original state when PRBS is disabled.
- On MX10003 routers, the PRBS58 pattern is supported only on JNP-MIC1-MACSEC MIC.

Issue the "[prbs-test-stop](#)" on page 1242 command to stop collecting the PRBS statistics .

For the step-by-step procedure on how to collect and view the PRBS statistics, refer "[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#)" on page 681.

## Options

***ifd-name***—Name of the interface.

***type***—Specifies the pattern type, that is in the range from 7 to 58.

Following pattern types are supported:

| Pattern Type | Pattern Name                                                                          |
|--------------|---------------------------------------------------------------------------------------|
| 7            | PRBS7                                                                                 |
| 9            | PRBS9                                                                                 |
| 11           | PRBS11                                                                                |
| 15           | PRBS15                                                                                |
| 23           | PRBS23                                                                                |
| 31           | PRBS31<br><b>NOTE:</b> Recommended pattern to check the quality of the received link. |
| 58           | PRBS58                                                                                |

**direction**—Specifies to configure transmit or receive PRBS pattern.

**flip**—Specifies if the pattern bits needs to be flipped or not.

## Required Privilege Level

view

## Release Information

Command introduced in Junos OS Release 19.2R1.

## RELATED DOCUMENTATION

[prbs-test-stop | 1242](#)

[show interfaces prbs-stats | 1703](#)

*clear interfaces statistics*

## prbs-test-stop

### IN THIS SECTION

- Syntax | 1242
- Description | 1242
- Options | 1242
- Required Privilege Level | 1243
- Sample Output | 1243
- Release Information | 1243

### Syntax

```
test interfaces ifd-name prbs-test-stop direction (0|1)
```

### Description

Use this command to stop collecting the Pseudo Random Binary Sequence (PRBS) statistics that is initiated using "[prbs-test-start](#)" on [page 1239](#) command. 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.

For the step-by-step procedure on how to collect and view the PRBS statistics, refer "[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#)" on [page 681](#).

### Options

- |                  |                                                |
|------------------|------------------------------------------------|
| <i>ifd-name</i>  | Name of the interface.                         |
| <i>direction</i> | Direction to transmit or receive PRBS pattern. |

## Required Privilege Level

view

## Sample Output

### Stopping PRBS test statistics collection

```
user@host> test interface et-0/1/2 prbs-test-stop direction 1
```

## Release Information

Command introduced in Junos OS Release 19.2R1.

## RELATED DOCUMENTATION

---

[prbs-test-start](#) | [1239](#)

---

[show interfaces prbs-stats](#) | [1703](#)

---

*clear interfaces statistics*

---

[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#) | [681](#)

## request interface link-degrade-recover

### IN THIS SECTION

- [Syntax](#) | [1244](#)
- [Description](#) | [1244](#)
- [Options](#) | [1244](#)
- [Required Privilege Level](#) | [1244](#)
- [Output Fields](#) | [1244](#)
- [Sample Output](#) | [1244](#)
- [Release Information](#) | [1247](#)

## Syntax

```
request interface link-degrade-recover interfaces-name
```

## Description

Manually recover a degraded physical link. Manual recovery is used when the interface has any Layer 2 and Layer 3 protocols that prevents autorecovery. This command is applicable only if you have configured the manual link recovery option on the interface.

**NOTE:** Manual recovery option is recommended for user deployments that have static route configurations causing the remote end of the link to start forwarding packets (as soon as the physical link is up) while auto-recovery is in progress.

## Options

*interfaces-name* Name of the interface.

## Required Privilege Level

View

## Output Fields

When you enter this command, Junos OS displays the status of your request.

## Sample Output

### Manual recovery

```
user@host>run request interface link-degrade-recover xe-9/1/11
FPC 9 PIC 1 PORT 11 Link Degrade Recovery Started
```

## Interface status when link degrade is enabled

```

user@host>run show interfaces xe-9/1/11
Physical interface: xe-9/1/11, Enabled, Physical link is Up
Interface index: 181, SNMP ifIndex: 664
 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: None, Source filtering: Disabled,
 Flow control: Enabled, Speed Configuration: Auto
 Pad to minimum frame size: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Schedulers : 0
 Current address: 28:8a:1c:c9:0e:32, Hardware address: 28:8a:1c:c9:0e:32
 Last flapped : 2017-10-25 01:53:17 PDT (00:00:10 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects: None
 PCS statistics
 Bit errors Seconds
 Errored blocks 0
 Link Degrade :
 Link Monitoring : Enable
 Link Degrade Set Threshold : 1E-8
 Link Degrade Clear Threshold : 1E-11
 Link Degrade War Set Threshold : 1E-9
 Link Degrade War Clear Threshold : 1E-10
 Estimated BER : <= 1E-16
 Link-degrade event : Seconds
 Count State
 0 OK
 Interface transmit statistics: Disabled

Logical interface xe-9/1/11.0 (Index 32368) (SNMP ifIndex 33153)
 Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol inet, MTU: 1500

```

```

Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold
cnt: 0, NH drop cnt: 0
 Flags: Sendbcst-pkt-to-re
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
 Protocol multiservice, MTU: Unlimited
 Flags: Is-Primary

```

## Interface status when the defect is active

```

user@host>run show interfaces xe-9/1/11
Physical interface: xe-9/1/11, Enabled, Physical link is Down
Interface index: 181, SNMP ifIndex: 664
 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: None, Source filtering: Disabled,
 Flow control: Enabled, Speed Configuration: Auto
 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 : 0
 Current address: 28:8a:1c:c9:0e:32, Hardware address: 28:8a:1c:c9:0e:32
 Last flapped : 2017-10-25 01:54:09 PDT (00:00:03 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : LINK
 Active defects: LINK, LOCAL-FAULT
 PCS statistics
 Bit errors Seconds
 Errored blocks 0
 Link Degrade :
 Link Monitoring : Enable
 Link Degrade Set Threshold : 1E-8
 Link Degrade Clear Threshold : 1E-11
 Link Degrade War Set Threshold : 1E-9
 Link Degrade War Clear Threshold : 1E-10
 Estimated BER : 1E-4
 Link-degrade event : Seconds
 Count State

```

```

4
1 Defect Active
Interface transmit statistics: Disabled

Logical interface xe-9/1/11.0 (Index 32368) (SNMP ifIndex 33153)
 Flags: Device-Down SNMP-Traps 0x4004000 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol inet, MTU: 1500
 Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold
cnt: 0, NH drop cnt: 0
 Flags: Sendbcast-pkt-to-re
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255
 Protocol multiservice, MTU: Unlimited
 Flags: Is-Primary

```

## Release Information

Command introduced in Junos OS Release 15.1.

## RELATED DOCUMENTATION

[Link Degradation Monitoring Overview | 600](#)

[link-degrade-monitor | 955](#)

[thresholds | 1179](#)

[recovery | 1101](#)

## request interface mc-ae switchover (Multichassis Link Aggregation)

### IN THIS SECTION

- [Syntax | 1248](#)
- [Description | 1248](#)
- [Options | 1248](#)

- [Required Privilege Level | 1248](#)
- [Output Fields | 1249](#)
- [Sample Output | 1249](#)
- [Sample Output | 1249](#)
- [Release Information | 1249](#)

## Syntax

```
request interface mc-ae switchover
<immediate> mcae-id mcae-id;
mcae-id mcae-id;
```

## Description

Manually revert egress traffic from the active node to the designated preferred node of a multichassis aggregated Ethernet interface. You can use this command to manually switch over traffic to the preferred node when the **switchover-mode** statement for the multichassis aggregated Ethernet interface is configured as **non-revertive** at the **[edit interfaces aeX mc-ae]** hierarchy level.

**NOTE:** To run this command successfully, the **status-control** statement should be configured as **active** at the **[edit interfaces aeX mc-ae]** hierarchy level.

## Options

**immediate** (Optional) Trigger immediate switchover to the preferred node. If this option is not configured, Junos OS waits for the timer configured using the **revert-time** statement at the **[edit interfaces aeX mc-ae]** hierarchy level to expire before it triggers the switchover.

**mcae-id *mcae-id*** Triggers switchover for the specified mc-ae interface.

## Required Privilege Level

view

## Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**request interface mc-ae switchover immediate mcae-id**

```
user@host >request interface mc-ae switchover immediate mcae-id 2
MCAE: Switchover Done
```

## Sample Output

**request interface mc-ae switchover mcae-id**

```
user@host >request interface mc-ae switchover mcae-id 2
Switchover In Progress: Please check after 1 minutes,
Use "show interfaces mc-ae revertive-info" to check for the status
```

## Release Information

Command introduced in Junos OS Release 13.3.

## RELATED DOCUMENTATION

[Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up](#)

[Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on MX Series Routers](#)

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

### IN THIS SECTION

- Syntax | 1250
- Description | 1250
- Options | 1251
- Required Privilege Level | 1251
- Output Fields | 1251
- Sample Output | 1251
- Release Information | 1251

### Syntax

```
request interface (revert | switchover) aex
```

### Description

Manually revert egress traffic from the designated backup link to the designated primary link of an aggregated Ethernet interface for which link protection is enabled, or manually switch egress traffic from the primary link to the backup link. This traffic includes transit traffic and local traffic originated on the router itself.

**NOTE:** When link protection is enabled on an aggregated Ethernet interface, if the primary link fails, the router automatically routes egress traffic to the backup link. However, the router does not automatically route egress traffic back to the primary link when the primary link is subsequently reestablished. Instead, you manually control when to have traffic diverted back to the primary link by issuing the **request interface (revert | switchover)** (Aggregated Ethernet Link Protection) operational command and specifying the **revert** keyword.

On M Series and T Series routers, use the **request interface (revert | switchover)** (Adaptive Services) operational command to manually revert to the primary adaptive services interface or link services

interface, or to switch from the primary to the secondary interface. For information about this command, see *request interface (revert | switchover) (Adaptive Services)*.

## Options

|                   |                                                                     |
|-------------------|---------------------------------------------------------------------|
| <b>revert</b>     | Restores egress traffic processing to the primary link.             |
| <b>switchover</b> | Transfers egress traffic processing to the secondary (backup) link. |
| <b>aeX</b>        | Aggregated Ethernet logical interface number: 0 through 15.         |

## Required Privilege Level

view

## Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request interface revert

```
user@host >request interface revert ae1
```

## Release Information

Command introduced in Junos OS Release 8.3.

## request lacp link-switchover

### IN THIS SECTION

- [Syntax | 1252](#)
- [Description | 1252](#)

- [Options | 1252](#)
- [Required Privilege Level | 1252](#)
- [Output Fields | 1252](#)
- [Sample Output | 1253](#)
- [Release Information | 1253](#)

## Syntax

```
request lacp link-switchover aex
```

## Description

Manually switch aggregated Ethernet active or standby LACP links.

**NOTE:** Because this command overrides LACP priority calculations, we strongly recommend that you use this command only when the actor (in this case, the Juniper Networks router) is controlling the active or standby link and the partner (peer) is following. This scenario occurs when you configure only the actor for link protection.

## Options

**aex** Aggregated Ethernet logical interface number: 0 through 15.

## Required Privilege Level

view

## Output Fields

When you enter this command, you are provided feedback on the status of your request. To view the switchover, use the **show lacp interfaces** command.

## Sample Output

**request lacp link-switchover aeX**

```
user@host >request lacp link-switchover ae0ae0: Request succeeded
```

## Release Information

Command introduced in Junos OS Release 9.3.

## show chassis hardware

### IN THIS SECTION

- [Syntax | 1254](#)
- [Syntax \(EX Series, MX104, MX204, MX2010, MX2020, MX10003, MX10008, and MX2008 Universal Routing Platforms\) | 1254](#)
- [Syntax \(TX Matrix Router\) | 1254](#)
- [Syntax \(TX Matrix Plus Router\) | 1254](#)
- [Syntax \(MX Series Routers\) | 1255](#)
- [Syntax \(QFX Series\) | 1255](#)
- [Description | 1255](#)
- [Options | 1256](#)
- [Additional Information | 1257](#)
- [Required Privilege Level | 1258](#)
- [Output Fields | 1258](#)
- [Sample Output | 1262](#)
- [Release Information | 1274](#)

## Syntax

```
show chassis hardware
<detail | extensive>
<clei-models>
<models>
```

## Syntax (EX Series, MX104, MX204, MX2010, MX2020, MX10003, MX10008, and MX2008 Universal Routing Platforms)

```
show chassis hardware
<clei-models>
<detail | extensive>
<models>
<satellite [slot-id slot-id | device-alias alias-name]>
```

## Syntax (TX Matrix Router)

```
show chassis hardware
<clei-models>
<detail | extensive>
<models>
<lcc number | scc>
```

## Syntax (TX Matrix Plus Router)

```
show chassis hardware
<clei-models>
<detail | extensive>
<models>
<lcc number | sfc number>
```

## Syntax (MX Series Routers)

```
show chassis hardware
<detail | extensive>
<clei-models>
<models>
<all-members>
<local>
<member member-id>
```

## Syntax (QFX Series)

```
show chassis hardware
<detail | extensive>
<clei-models>
<interconnect-device name>
<node-device name>
<models>
```

## Description

Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number.

In the EX Series switch command output, FPC refers to the following:

- On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC *number* is always 0.
- On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC *number* equals the member ID, from 0 through 9.
- On EX8208 and EX8216 switches—Refers to a line card; FPC *number* equals the slot number for the line card.

On QFX3500, QFX5100, and OCX Series standalone switches, and PTX1000 routers both the FPC and FPC *number* are always 0.

On T4000 Type 5 FPCs, there are no **top temperature sensor** or **bottom temperature sensor** parameters. Instead, **fan intake temperature sensor** and **fan exhaust temperature sensors** parameters are displayed.

Starting from Junos OS Release 11.4, the output of the **show chassis hardware models** operational mode command displays the enhanced midplanes FRU model numbers (CHAS-BP3-MX240-S, CHAS-BP3-MX480-S or CHAS-BP3-MX960-S) based on the router. Prior to release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through the Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 14.1, the output of the **show chassis hardware detail | extensive | clei-models | models** operational mode command displays the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-P1A) and other components in a PTX5000 Packet Transport Router.

## Options

|                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>none</b>                            | Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>clei-models</b>                     | (Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>detail</b>                          | (Optional) Include RAM and disk information in output.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>extensive</b>                       | (Optional) Display ID EEPROM information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>all-members</b>                     | (MX Series routers only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>interconnect-device <i>name</i></b> | (QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>lcc <i>number</i></b>               | (TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.<br><br>Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> </ul> |

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

|                                                                                    |                                                                                                                                                                                                                   |
|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>local</b>                                                                       | (MX Series routers only) (Optional) Display hardware-specific information for the local Virtual Chassis members.                                                                                                  |
| <b>member <i>member-id</i></b>                                                     | (MX Series routers and EX Series switches) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> variable with a value 0 or 1. |
| <b>models</b>                                                                      | (Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.                                                                             |
| <b>node-device <i>name</i></b>                                                     | (QFabric systems only) (Optional) Display hardware-specific information for the Node device.                                                                                                                      |
| <b>satellite [slot-id<br/><i>slot-id</i>   device-alias<br/><i>alias-name</i>]</b> | (Junos Fusion only) (Optional) Display hardware information for the specified satellite device in a Junos Fusion, or for all satellite devices in the Junos Fusion if no satellite devices are specified.         |
| <b>scc</b>                                                                         | (TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (switch-card chassis).                                                                                                   |
| <b>sfc <i>number</i></b>                                                           | (TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (switch-fabric chassis). Replace <i>number</i> variable with 0.                                                |

## Additional Information

The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines, as shown in [Table 136 on page 1257](#).

**Table 136: Routing Engines Displaying DIMM Information**

| Routing Engines             | Routers                         |
|-----------------------------|---------------------------------|
| RE-S-1800x2 and RE-S-1800x4 | MX240, MX480, and MX960 routers |
| RE-A-1800x2                 | M120 and M320 routers           |

In Junos OS Release 11.4 and later, the output for the **show chassis hardware models** operational mode command for MX Series routers display the enhanced midplanes FRU model numbers—CHAS-BP3-MX240-S, CHAS-BP3-MX480-S, or CHAS-BP3-MX960-S—based on the router. In releases before Junos OS Release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 17.3R1, the output of the **show chassis hardware** command displays the mode in which vMX is running (performance mode or lite mode) in the part number field for the FPC. **RIOT-PERF** indicates performance mode and **RIOT-LITE** indicates lite mode.

## Required Privilege Level

view

## Output Fields

[Table 137 on page 1258](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

**Table 137: show chassis hardware Output Fields**

| Field Name                    | Field Description                                                                                                                                                                                                                                 | Level of Output  |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Item</b>                   | Show information about the device hardware.                                                                                                                                                                                                       | All levels       |
| <b>Version</b>                | Revision level of the chassis component.                                                                                                                                                                                                          | All levels       |
| <b>Part number</b>            | Part number of the chassis component.                                                                                                                                                                                                             | All levels       |
| <b>Serial number</b>          | Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis. | All levels       |
| <b>Assb ID or Assembly ID</b> | ( <b>extensive</b> keyword only) Identification number that describes the FRU hardware.                                                                                                                                                           | <b>extensive</b> |

Table 137: show chassis hardware Output Fields *(Continued)*

| Field Name              | Field Description                                                                                                                                                                                                                                              | Level of Output  |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Assembly Version</b> | ( <b>extensive</b> keyword only) Version number of the FRU hardware.                                                                                                                                                                                           | <b>extensive</b> |
| <b>Assembly Flags</b>   | ( <b>extensive</b> keyword only) Flags.                                                                                                                                                                                                                        | <b>extensive</b> |
| <b>FRU model number</b> | ( <b>clei-models</b> , <b>extensive</b> , and <b>models</b> keyword only) Model number of the FRU hardware component.                                                                                                                                          | none specified   |
| <b>CLEI code</b>        | ( <b>clei-models</b> and <b>extensive</b> keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1. | none specified   |
| <b>EEPROM Version</b>   | ID EEPROM version used by the hardware component: <b>0x00</b> (version 0), <b>0x01</b> (version 1), or <b>0x02</b> (version 2).                                                                                                                                | <b>extensive</b> |

Table 137: show chassis hardware Output Fields (Continued)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Description</b> | <p data-bbox="444 369 894 401">Brief description of the hardware item:</p> <ul data-bbox="444 432 1195 1713" style="list-style-type: none"> <li data-bbox="444 432 737 464">• Type of power supply.</li> <li data-bbox="444 495 1195 569">• Type of PIC. If the PIC type is not supported on the current software release, the output states <b>Hardware Not Supported</b>.</li> <li data-bbox="444 600 1073 674">• Type of FPC: <b>FPC Type 1, FPC Type 2, FPC Type 3, FPC Type 4</b> , or <b>FPC TypeOC192</b>.</li> </ul> <p data-bbox="480 705 1101 737">On EX Series switches, a brief description of the FPC.</p> <p data-bbox="480 768 1170 842">The following list shows the PIM abbreviation in the output and the corresponding PIM name.</p> <ul data-bbox="480 873 1170 1713" style="list-style-type: none"> <li data-bbox="480 873 1170 947">• <b>2x FE</b>—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM</li> <li data-bbox="480 978 911 1010">• <b>4x FE</b>—4-port Fast Ethernet ePIM</li> <li data-bbox="480 1041 1162 1115">• <b>1x GE Copper</b>—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port)</li> <li data-bbox="480 1146 1154 1178">• <b>1x GE SFP</b>—SFP Gigabit Ethernet ePIM (one fiber port)</li> <li data-bbox="480 1209 878 1241">• <b>2x Serial</b>—Dual-port serial PIM</li> <li data-bbox="480 1272 813 1304">• <b>2x T1</b>—Dual-port T1 PIM</li> <li data-bbox="480 1335 813 1367">• <b>2x E1</b>—Dual-port E1 PIM</li> <li data-bbox="480 1398 1049 1430">• <b>2x CT1E1</b>—Dual-port channelized T1/E1 PIM</li> <li data-bbox="480 1461 813 1493">• <b>1x T3</b>—T3 PIM (one port)</li> <li data-bbox="480 1524 813 1556">• <b>1x E3</b>—E3 PIM (one port)</li> <li data-bbox="480 1587 959 1619">• <b>4x BRI S/T</b>—4-port ISDN BRI S/T PIM</li> <li data-bbox="480 1650 919 1682">• <b>4x BRI U</b>—4-port ISDN BRI U PIM</li> </ul> | All levels      |

Table 137: show chassis hardware Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <li>• <b>1x ADSL Annex A</b>—ADSL 2/2+ Annex A PIM (one port, for POTS)</li> <li>• <b>1x ADSL Annex B</b>—ADSL 2/2+ Annex B PIM (one port, for ISDN)</li> <li>• <b>2x SHDSL (ATM)</b>—G SHDSL PIM (2-port two-wire module or 1-port four-wire module)</li> <li>• <b>1x TGM550</b>—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog <b>LINE</b> ports, and two analog <b>TRUNK</b> ports)</li> <li>• <b>1x DS1 TIM510</b>—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup)</li> <li>• <b>4x FXS, 4xFX0, TIM514</b>—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog <b>LINE</b> ports and four analog <b>TRUNK</b> ports)</li> <li>• <b>4x BRI TIM521</b>—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports)</li> <li>• <b>Crypto Accelerator Module</b>—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services</li> <li>• <b>MPC M 16x 10GE</b>—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.)</li> <li>• For hosts, the Routing Engine type.</li> <li>• For small form-factor pluggable transceiver (SFP) modules, the type of fiber: <b>LX, SX, LH</b>, or <b>T</b>.</li> <li>• LCD description for EX Series switches (except EX2200 switches).</li> </ul> |                 |

Table 137: show chassis hardware Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <li>• <b>MPC2</b>—1-port MPC2 that supports two separate slots for MICs.</li> <li>• <b>MPC3E</b>—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs.</li> <li>• 100GBASE-LR4, pluggable CFP optics</li> <li>• Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy.</li> <li>• Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs).</li> <li>• <b>MPC4E</b>—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE on MX2020, MX960, MX480, and MX240 routers.</li> <li>• LCD description for MX Series routers</li> </ul> |                 |

## Sample Output

### show chassis hardware (MX10008 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

```
Item Version Part number Serial number Description
Chassis DE538 JNP10008 [MX10008]
Midplane REV 27 750-054097 ACPD6954 Midplane 8
Routing Engine 0 BUILTIN BUILTIN RE X10
Routing Engine 1 BUILTIN BUILTIN RE X10 128
CB 0 REV 10 750-079562 CAKF2158 Control Board
```

|            |        |            |                |                       |
|------------|--------|------------|----------------|-----------------------|
| CB 1       | REV 05 | 711-065897 | CAJG2680       | Control Board         |
| FPC 1      | REV 04 | 750-084779 | CAKN5706       | JNP10K-LC2101         |
| CPU        | REV 05 | 750-073391 | CAKJ2864       | LC 2101 PMB           |
| PIC 0      |        | BUILTIN    | BUILTIN        | 4xQSFP28 SYNCE        |
| Xcvr 0     | REV 01 | 740-046565 | XXL0BQM        | QSFP+-40G-SR4         |
| Xcvr 1     | REV 01 | 740-032986 | QB350242       | QSFP+-40G-SR4         |
| Xcvr 2     | REV 01 | 740-054053 | QE408285       | QSFP+-4X10G-SR        |
| Xcvr 3     | REV 01 | 740-046565 | QF3300Z9       | QSFP+-40G-SR4         |
| PIC 1      |        | BUILTIN    | BUILTIN        | 4xQSFP28 SYNCE        |
| Xcvr 0     | REV 01 | 740-067442 | QJ2200LD       | QSFP+-40G-SR4         |
| Xcvr 1     | REV 01 | 740-038153 | APF170500382DP | QSFP+-40G-CU3M        |
| Xcvr 2     | REV 01 | 740-067442 | QI4302LC       | QSFP+-40G-SR4         |
| PIC 2      |        | BUILTIN    | BUILTIN        | 4xQSFP28 SYNCE        |
| Xcvr 0     | REV 01 | 740-067442 | 1ACP1335119    | QSFP+-40G-SR4         |
| Xcvr 1     | REV 01 | 740-067442 | 1ACP1313156    | QSFP+-40G-SR4         |
| Xcvr 2     | REV 01 | 740-067442 | QK050040       | QSFP+-40G-SR4         |
| Xcvr 3     | REV 01 | 740-067442 | QJ2201BG       | QSFP+-40G-SR4         |
| PIC 3      |        | BUILTIN    | BUILTIN        | 4xQSFP28 SYNCE        |
| Xcvr 0     | REV 01 | 740-058734 | 1ECQ12400CS    | QSFP-100GBASE-SR4     |
| Xcvr 1     | REV 01 | 740-046565 | QF3300ZX       | QSFP+-40G-SR4         |
| Xcvr 2     | REV 01 | 740-061405 | 1ECQ12510FH    | QSFP-100G-SR4-T2      |
| Xcvr 3     | REV 01 | 740-032986 | QB491182       | QSFP+-40G-SR4         |
| PIC 4      |        | BUILTIN    | BUILTIN        | 4xQSFP28 SYNCE        |
| Xcvr 0     | REV 01 | 740-067442 | QJ2200D5       | QSFP+-40G-SR4         |
| Xcvr 1     | REV 01 | 740-054053 | XXS0L95        | QSFP+-4X10G-SR        |
| PIC 5      |        | BUILTIN    | BUILTIN        | 4xQSFP28 SYNCE        |
| Xcvr 0     | REV 01 | 740-054053 | QE251550       | QSFP+-4X10G-SR        |
| Xcvr 1     | REV 01 | 740-054053 | XZB01D5        | QSFP+-4X10G-SR        |
| Xcvr 3     | REV 01 | 740-046565 | QI1402F9       | QSFP+-40G-SR4         |
| FPD Board  | REV 07 | 711-054687 | ACPF2896       | Front Panel Display   |
| PEM 1      | REV 02 | 740-049388 | 1EDL62102PR    | Power Supply AC       |
| PEM 2      | REV 02 | 740-049388 | 1EDL60300H2    | Power Supply AC       |
| PEM 4      | REV 02 | 740-049388 | 1EDL603003Z    | Power Supply AC       |
| PEM 5      | REV 01 | 740-049388 | 1EDL339001B    | Power Supply AC       |
| FTC 0      | REV 14 | 750-050108 | ACNW3344       | Fan Controller 8      |
| FTC 1      | REV 14 | 750-050108 | ACPE3978       | Fan Controller 8      |
| Fan Tray 0 | REV 09 | 760-054372 | ACNV5507       | Fan Tray 8            |
| Fan Tray 1 | REV 09 | 760-054372 | ACNV5371       | Fan Tray 8            |
| SFB 0      | REV 25 | 750-050058 | ACPH6821       | Switch Fabric (SIB) 8 |
| SFB 1      |        |            |                |                       |
| SFB 2      | REV 24 | 750-050058 | ACNZ0641       | Switch Fabric (SIB) 8 |

|       |        |            |          |                       |
|-------|--------|------------|----------|-----------------------|
| SFB 3 | REV 27 | 750-050058 | ACPH9127 | Switch Fabric (SIB) 8 |
| SFB 5 | REV 24 | 750-050058 | ACNX7396 | Switch Fabric (SIB) 8 |

### show chassis hardware (PTX10008 Router)

```

root@host> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis Version Part number Serial number JNP10008 [PTX10008]
Midplane REV 28 750-054097 ACP22394 Midplane 8
Routing Engine 0 BUILTIN BUILTIN RE-PTX-2X00x4
Routing Engine 1 BUILTIN BUILTIN RE-PTX-2X00x4
CB 0 REV 04 750-068820 ACPT5303 Control Board
CB 1 REV 04 750-068820 ACPR1627 Control Board
FPC 6 REV 13 750-068822 ACPB2753 LC1102 - 12C / 36Q /
14
 4X
CPU BUILTIN BUILTIN FPC CPU
PIC 0 BUILTIN BUILTIN 12x100GE/36x40GE/
144x10
 GE
Xcvr 0 REV 01 740-067442 XV304N6 QSFP+-40G-SR4
Xcvr 1 REV 01 740-067442 XV30A5M QSFP+-40G-SR4
Xcvr 2 REV 01 740-067442 XV300HC QSFP+-40G-SR4
Xcvr 3 REV 01 740-067443 XU20L17 QSFP+-40G-SR4
Xcvr 4 REV 01 740-067442 XV303XG QSFP+-40G-SR4
Xcvr 5 REV 01 740-067443 XV306QC QSFP+-40G-SR4
Xcvr 6 REV 01 740-067442 XV303Y7 QSFP+-40G-SR4
Xcvr 7 REV 01 740-067443 XX60DMR QSFP+-40G-SR4
Xcvr 9 REV 01 740-067443 XX60DNY QSFP+-40G-SR4
Xcvr 10 REV 01 740-054053 QF4605WF QSFP+-4X10G-SR
Xcvr 13 REV 01 740-058734 1ECQ115007D QSFP-100GBASE-SR4
Xcvr 15 REV 01 740-046565 QH06035R QSFP+-40G-SR4
Xcvr 16 REV 01 740-046565 QH0602KC QSFP+-40G-SR4
Xcvr 17 REV 01 740-046565 QH0507PA QSFP+-40G-SR4
Xcvr 18 REV 01 740-046565 QH06035M QSFP+-40G-SR4
Xcvr 24 REV 01 740-046565 QH0507QL QSFP+-40G-SR4
Xcvr 25 REV 01 740-067443 XV20CWP QSFP+-40G-SR4
Xcvr 34 REV 01 740-046565 QH06035U QSFP+-40G-SR4
Xcvr 35 REV 01 740-067443 XX60DN9 QSFP+-40G-SR4
FPC 7 REV 41 750-051357 ACPL3446 LC1101 - 30C / 30Q / 96X

```

| CPU            |        | BUILTIN    | BUILTIN     | FPC CPU                  |
|----------------|--------|------------|-------------|--------------------------|
| PIC 0          |        | BUILTIN    | BUILTIN     | 30x100GE/30x40GE/96x10GE |
| Xcvr 0         | REV 01 | 740-067443 | XX60DPC     | QSFP+-40G-SR4            |
| Xcvr 1         | REV 01 | 740-054053 | QF4605W7    | QSFP+-4X10G-SR           |
| Xcvr 2         | REV 01 | 740-067443 | XX60DP8     | QSFP+-40G-SR4            |
| Xcvr 3         | REV 01 | 740-067442 | XV30FYM     | QSFP+-40G-SR4            |
| Xcvr 4         | REV 01 | 740-067442 | 1ACP133406Z | QSFP+-40G-SR4            |
| Xcvr 5         | REV 01 | 740-067443 | XX60DP5     | QSFP+-40G-SR4            |
| Xcvr 8         | REV 01 | 740-046565 | QH060355    | QSFP+-40G-SR4            |
| Xcvr 12        | REV 01 | 740-058734 | 1ECQ115008C | QSFP-100GBASE-SR4        |
| Xcvr 15        | REV 01 | 740-046565 | QH0602KG    | QSFP+-40G-SR4            |
| Xcvr 16        | REV 01 | 740-046565 | QH0602LG    | QSFP+-40G-SR4            |
| Xcvr 17        | REV 01 | 740-046565 | QH06035S    | QSFP+-40G-SR4            |
| Xcvr 18        | REV 01 | 740-046565 | QH0602KS    | QSFP+-40G-SR4            |
| Xcvr 24        | REV 01 | 740-067443 | QI2902DP    | QSFP+-40G-SR4            |
| Power Supply 2 | REV 02 | 740-049388 | 1EDL70200NP | Power Supply AC          |
| Power Supply 3 | REV 02 | 740-049388 | 1EDL603005C | Power Supply AC          |
| Power Supply 4 | REV 02 | 740-049388 | 1EDL70200P1 | Power Supply AC          |
| Power Supply 5 | REV 02 | 740-049388 | 1EDL70200B7 | Power Supply AC          |
| FTC 0          | REV 16 | 750-050108 | ACPK8682    | Fan Controller 8         |
| FTC 1          | REV 16 | 750-050108 | ACPR9530    | Fan Controller 8         |
| Fan Tray 0     | REV 10 | 760-054372 | ACPR9509    | Fan Tray 8               |
| Fan Tray 1     | REV 10 | 760-054372 | ACPV7260    | Fan Tray 8               |
| SIB 1          | REV 28 | 750-050058 | ACPV6306    | Switch Fabric 8          |
| SIB 2          | REV 28 |            | ACPR2569    | Switch Fabric 8          |
| SIB 3          | REV 28 | 750-05     | ACPW7402    | Switch Fabric 8          |
| SIB 4          | REV 28 | 750-050058 | ACPR2577    | Switch Fabric 8          |
| FPD Board      | REV 07 | 711-054687 | ACPM4965    | Front Panel Display      |

### show chassis hardware clei-models (PTX10016 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item Version Part number CLEI code FRU model number
Midplane REV 24 750-077138 CMMUN00ARA JNP10016
CB 0 REV 04 711-065897 PROTOXCLEI PROTO-ASSEMBLY
CB 1 REV 05 711-065897 PROTOXCLEI PROTO-ASSEMBLY
FPC 2
 PIC 0 BUILTIN
FPC 4 REV 35 750-071976 CMUIANABAA JNP10K-LC1101
 PIC 0 BUILTIN

```

|                 |        |                       |            |                 |
|-----------------|--------|-----------------------|------------|-----------------|
| FPC 5<br>PIC 0  | REV 13 | 750-068822<br>BUILTIN | CMUIAM9BAC | QFX10000-36Q    |
| FPC 6<br>PIC 0  | REV 41 | 750-071976<br>BUILTIN | CMUIANABAB | JNP10K-LC1101   |
| FPC 7<br>PIC 0  | REV 35 | 750-071976<br>BUILTIN | CMUIANABAA | JNP10K-LC1101   |
| FPC 8<br>PIC 0  | REV 35 | 750-071976<br>BUILTIN | CMUIANABAA | JNP10K-LC1101   |
| FPC 9<br>PIC 0  | REV 41 | 750-071976<br>BUILTIN | CMUIANABAB | JNP10K-LC1101   |
| FPC 10<br>PIC 0 | REV 35 | 750-071976<br>BUILTIN | CMUIANABAA | JNP10K-LC1101   |
| FPC 11<br>PIC 0 | REV 35 | 750-071976<br>BUILTIN | CMUIANABAA | JNP10K-LC1101   |
| FPC 13<br>PIC 0 | REV 41 | 750-071976<br>BUILTIN | CMUIANABAB | JNP10K-LC1101   |
| FPC 15<br>PIC 0 | REV 37 | 750-071976<br>BUILTIN | CMUIANABAA | JNP10K-LC1101   |
| Power Supply 0  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 1  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 2  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 3  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 4  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 5  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 6  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 7  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 8  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Power Supply 9  | REV 01 | 740-073147            | CMUPADPBAA | JNP10K-PWR-DC   |
| Fan Tray 0      |        |                       |            | QFX5100-FAN-AFO |
| Fan Tray 1      |        |                       |            | QFX5100-FAN-AFO |
| SIB 0           | REV 15 | 750-077140            | CMUCAH6CAA | JNP10016-SF     |
| SIB 1           | REV 15 | 750-077140            | CMUCAH6CAA | JNP10016-SF     |
| SIB 2           | REV 15 | 750-077140            | CMUCAH6CAA | JNP10016-SF     |
| SIB 3           | REV 15 | 750-077140            | CMUCAH6CAA | JNP10016-SF     |
| SIB 4           | REV 15 | 750-077140            | CMUCAH6CAA | JNP10016-SF     |
| SIB 5           | REV 15 | 750-077140            | CMUCAH6CAA | JNP10016-SF     |
| FPD Board       | REV 07 | 711-054687            |            |                 |

### show chassis hardware detail (EX9251 Switch)

```

user@switch> show chassis hardware
Hardware inventory:
Item Version Part number Serial number Description
Chassis BLANK EX9251
Routing Engine 0 BUILTIN BUILTIN RE-S-2X00x6
CB 0 REV 05 750-069579 CAGT1382 EX9251
FPC 0 BUILTIN BUILTIN MPC
 PIC 0 BUILTIN BUILTIN 4XSFP28 PIC
 Xcvr 0 REV 01 740-044512 APF14500007NHC QSFP+-40G-CU50CM
 Xcvr 2 REV 01 740-046565 QH21035H QSFP+-40G-SR4
 PIC 1 BUILTIN BUILTIN 8XSFP PIC
 Xcvr 0 REV 01 740-031980 AA15393URH7 SFP+-10G-SR
 Xcvr 1 REV 01 740-031980 AA162832LVG SFP+-10G-SR
 Xcvr 2 REV 01 740-031980 MXA0NKJ SFP+-10G-SR
 Xcvr 3 REV 01 740-031980 MXA0K75 SFP+-10G-SR
 Xcvr 4 REV 01 740-021308 MXA138L SFP+-10G-SR
 Xcvr 5 REV 01 740-021308 13T5111102684 SFP+-10G-SR
 Xcvr 6 REV 01 740-021308 MXA138E SFP+-10G-SR
 Xcvr 7 REV 01 740-021308 MXA152N SFP+-10G-SR
PEM 0 REV 02 740-070749 1F186390060 AC AFO 650W PSU
PEM 1 REV 02 740-070749 1F186390045 AC AFO 650W PSU
Fan Tray 0 Fan Tray, Front to Back
Airflow - AFO
Fan Tray 1 Fan Tray, Front to Back
Airflow - AFO

```

### show chassis hardware extensive (T640 Router)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item Version Part number Serial number Description
Chassis T640
Jedec Code: 0x7fb0 EEPROM Version: 0x01
P/N: S/N:
Assembly ID: 0x0507 Assembly Version: 00.00
Date: 00-00-0000 Assembly Flags: 0x00
Version:
ID: Gibson LCC Chassis

```

Board Information Record:

Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

I2C Hex Data:

Address 0x00: 7f b0 01 ff 05 07 00 00 00 00 00 00 00 00 00 00

Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x20: ff 00 00 00 00

Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Midplane REV 04 710-002726 AX5633

Jedec Code: 0x7fb0 EEPROM Version: 0x01

P/N: 710-002726. S/N: AX5633.

Assembly ID: 0x0127 Assembly Version: 01.04

Date: 06-27-2001 Assembly Flags: 0x00

Version: REV 04.....

ID: Gibson Backplane

Board Information Record:

Address 0x00: ad 01 08 00 00 90 69 0e f8 00 ff ff ff ff ff ff

I2C Hex Data:

Address 0x00: 7f b0 01 ff 01 27 01 04 52 45 56 20 30 34 00 00

Address 0x10: 00 00 00 00 37 31 30 2d 30 30 32 37 32 36 00 00

Address 0x20: 53 2f 4e 20 41 58 35 36 33 33 00 00 00 1b 06 07

Address 0x30: d1 ff ff ff ad 01 08 00 00 90 69 0e f8 00 ff ff

Address 0x40: ff ff

FPM GBUS REV 02 710-002901 HE3245

...

FPM Display REV 02 710-002897 HA4873

...

CIP REV 05 710-002895 HA4729

...

PEM 1 RevX02 740-002595 MD21815 Power Entry Module

...

SCG 0 REV 04 710-003423 HF6023

...

SCG 1 REV 04 710-003423 HF6061

...

Routing Engine 0 REV 01 740-005022 210865700292 RE-3.0

...

CB 0 REV 06 710-002728 HE3614

...

FPC 1 REV 01 710-002385 HE3009 FPC Type 1

...

REV 06 710-001726 HC0010

**show chassis hardware interconnect-device (QFabric Systems)**

```

user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
Item Version Part number Serial number Description
Chassis REV 07
Midplane REV 07 750-021261 BH0208188289 QFX Midplane
CB 0 REV 07 750-021261 BH0208188289 QFXIC08-CB4S

```

**show chassis hardware lcc (TX Matrix Router)**

```

user@host> show chassis hardware lcc 0
lcc0-re0:

Hardware inventory:
Item Version Part number Serial number Description
Chassis REV 03
Midplane REV 03 710-005608 RA1408 T640 Backplane
FPM GBUS REV 09 710-002901 RA2784 T640 FPM Board
FPM Display REV 05 710-002897 RA2825 FPM Display
CIP REV 06 710-002895 HT0684 T Series CIP
PEM 0 Rev 11 740-002595 PM18483 Power Entry Module
PEM 1 Rev 11 740-002595 qb13984 Power Entry Module
SCG 0 REV 11 710-003423 HT0022 T640 Sonet Clock Gen.
Routing Engine 0 REV 13 740-005022 210865700363 RE-3.0 (RE-600)
CB 0 REV 03 710-007655 HW1195 Control Board (CB-T)
FPC 1 REV 05 710-007527 HM3245 FPC Type 2
 CPU REV 14 710-001726 HM1084 FPC CPU
 PIC 0 REV 02 750-007218 AZ1112 2x OC-12 ATM2 IQ, SMIR
 PIC 1 REV 02 750-007745 HG3462 4x OC-3 SONET, SMIR
 PIC 2 REV 14 750-001901 BA5390 4x OC-12 SONET, SMIR
 PIC 3 REV 09 750-008155 HS3012 2x G/E IQ, 1000 BASE
 SFP 0 NON-JNPR P1186TY SFP-S
 SFP 1 REV 01 740-007326 P11WLTF SFP-SX
MMB 1 REV 02 710-005555 HL7514 MMB-288mbit
PPB 0 REV 04 710-003758 HM4405 PPB Type 2
PPB 1 REV 04 710-003758 AV1960 PPB Type 2
FPC 2 REV 08 710-010154 HZ3578 E-FPC Type 3
 CPU REV 05 710-010169 HZ3219 FPC CPU-Enhanced
 PIC 0 REV 02 750-009567 HX2882 1x 10GE(LAN), XENPAK

```



|            |        |            |          |                    |
|------------|--------|------------|----------|--------------------|
| MIC 0      | REV 03 | 750-033307 | ZV6299   | MIC3-3D-10XGE-SFPP |
| MIC 1      | REV 03 | 750-033307 | ZV6268   | MIC3-3D-10XGE-SFPP |
| FPC 8      | REV 22 | 750-031089 | ZT9746   | MX-MPC2-3D         |
| MIC 0      | REV 26 | 750-028392 | ABBS1150 | MIC-3D-20GE-SFP    |
| MIC 1      | REV 26 | 750-028387 | ABBR9582 | MIC-3D-4XGE-XFP    |
| FPC 9      | REV 11 | 750-036284 | ZL3591   | MPCE-3D-16XGE-SFPP |
| ADC 0      | REV 05 | 750-043596 | CAAC2073 | 750-043596         |
| ADC 1      | REV 01 | 750-043596 | ZV4117   | 750-043596         |
| ADC 8      | REV 01 | 750-043596 | ZV4107   | 750-043596         |
| ADC 9      | REV 02 | 750-043596 | ZW1555   | 750-043596         |
| Fan Tray 0 | REV 2A | 760-046960 | ACAY0015 |                    |
| Fan Tray 1 | REV 2A | 760-046960 | ACAY0019 |                    |
| Fan Tray 2 | REV 2A | 760-046960 | ACAY0020 |                    |
| Fan Tray 3 | REV 2A | 760-046960 | ACAY0021 |                    |

### show chassis hardware node-device (QFabric Systems)

```

user@switch> show chassis hardware node-device node1
Routing Engine 0 BUILTIN BUILTIN QFX Routing Engine
node1 REV 05 711-032234 ED3694 QFX3500-48S4Q-AFI

CPU BUILTIN BUILTIN FPC CPU
PIC 0 BUILTIN BUILTIN 48x 10G-SFP+
Xcvr 8 REV 01 740-030658 AD0946A028B SFP+-10G-USR

```

### show chassis hardware scc (TX Matrix Router)

```

user@host> show chassis hardware scc
scc-re0:

Hardware inventory:
Item Version Part number Serial number Description
Chassis
Midplane REV 04 710-004396 RB0014 SCC Midplane
FPM GBUS REV 04 710-004617 HW9141 SCC FPM Board
FPM Display REV 04 710-004619 HS5950 SCC FPM
CIP 0 REV 01 710-010218 HV9151 SCC CIP
CIP 1 REV 01 710-010218 HV9152 SCC CIP
PEM 1 Rev 11 740-002595 QB13977 Power Entry Module

```

|                  |        |            |              |                       |
|------------------|--------|------------|--------------|-----------------------|
| Routing Engine 0 | REV 05 | 740-008883 | P11123900153 | RE-4.0 (RE-1600)      |
| CB 0             | REV 01 | 710-011709 | HR5964       | Control Board (CB-TX) |
| SPMB 0           | REV 09 | 710-003229 | HW5293       | T Series Switch CPU   |
| SIB 3            |        |            |              |                       |
| SIB 4            | REV 01 | 710-005839 | HW1177       | SIB-S8-F16            |
| B Board          | REV 01 | 710-005840 | HW1202       | SIB-S8-F16 (B)        |

### show chassis hardware sfc (TX Matrix Plus Router)

```

user@host> show chassis hardware sfc 0
sfc0-re0:

Hardware inventory:
Item Version Part number Serial number Description
Chassis REV 05 710-022574 TS4027 SFC Midplane
FPM Display REV 03 710-024027 DX0282 TXP FPM Display
CIP 0 REV 04 710-023792 DW4889 TXP CIP
CIP 1 REV 04 710-023792 DW4887 TXP CIP
PEM 0 Rev 07 740-027463 UM26368 Power Entry Module
Routing Engine 0 REV 01 740-026942 737A-1064 SFC RE
Routing Engine 1 REV 01 740-026942 737A-1082 SFC RE
CB 0 REV 09 710-022606 DW6099 SFC Control Board
CB 1 REV 09 710-022606 DW6096 SFC Control Board
SPMB 0 BUILTIN SFC Switch CPU
SPMB 1 BUILTIN SFC Switch CPU
SIB F13 0 REV 04 710-022600 DX0841 F13 SIB
 B Board REV 03 710-023431 DX0966 F13 SIB Mezz
SIB F13 1 REV 04 750-024564 DW5776 F13 SIB
 B Board REV 03 710-023431 DW9028 F13 SIB
SIB F13 3 REV 04 750-024564 DW5762 F13 SIB
 B Board REV 03 710-023431 DW9059 F13 SIB
SIB F13 4 REV 04 750-024564 DW5797 F13 SIB
 B Board REV 03 710-023431 DW9041 F13 SIB
SIB F13 6 REV 04 750-024564 DW5770 F13 SIB
 B Board REV 03 710-023431 DW9079 F13 SIB Mezz
SIB F13 7 REV 04 750-024564 DW5758 F13 SIB
 B Board REV 03 710-023431 DW9047 F13 SIB
SIB F13 8 REV 04 750-024564 DW5761 F13 SIB
 B Board REV 03 710-023431 DW9043 F13 SIB Mezz
SIB F13 9 REV 04 750-024564 DW5754 F13 SIB

```

|             |        |            |        |              |
|-------------|--------|------------|--------|--------------|
| B Board     | REV 03 | 710-023431 | DW9078 | F13 SIB Mezz |
| SIB F13 11  | REV 04 | 710-022600 | DX0826 | F13 SIB      |
| B Board     | REV 03 | 710-023431 | DX0967 | F13 SIB Mezz |
| SIB F13 12  | REV 04 | 750-024564 | DW5794 | F13 SIB      |
| B Board     | REV 03 | 710-023431 | DW9044 | F13 SIB Mezz |
| SIB F2S 0/0 | REV 05 | 710-022603 | DW7897 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7657 | NEO PMB      |
| SIB F2S 0/2 | REV 05 | 710-022603 | DW7833 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7526 | NEO PMB      |
| SIB F2S 0/4 | REV 05 | 710-022603 | DW7875 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7588 | NEO PMB      |
| SIB F2S 0/6 | REV 05 | 710-022603 | DW7860 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7589 | NEO PMB      |
| SIB F2S 1/0 | REV 04 | 710-022603 | DW4820 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW8510 | NEO PMB      |
| SIB F2S 1/2 | REV 05 | 710-022603 | DW7849 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7525 | NEO PMB      |
| SIB F2S 1/4 | REV 05 | 710-022603 | DW7927 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7556 | F2S SIB Mezz |
| SIB F2S 1/6 | REV 05 | 710-022603 | DW7866 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7651 | NEO PMB      |
| SIB F2S 2/0 | REV 05 | 710-022603 | DW7880 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7523 | NEO PMB      |
| SIB F2S 2/2 | REV 05 | 710-022603 | DW7895 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7591 | NEO PMB      |
| SIB F2S 2/4 | REV 05 | 710-022603 | DW7907 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7590 | NEO PMB      |
| SIB F2S 2/6 | REV 05 | 710-022603 | DW7785 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7524 | NEO PMB      |
| SIB F2S 3/0 | REV 05 | 710-022603 | DW7782 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7634 | NEO PMB      |
| SIB F2S 3/2 | REV 05 | 710-022603 | DW7793 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7548 | NEO PMB      |
| SIB F2S 3/4 | REV 05 | 710-022603 | DW7779 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7587 | NEO PMB      |
| SIB F2S 3/6 | REV 05 | 710-022603 | DW7930 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7505 | NEO PMB      |
| SIB F2S 4/0 | REV 05 | 710-022603 | DW7867 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7656 | NEO PMB      |
| SIB F2S 4/2 | REV 05 | 710-022603 | DW7917 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7640 | NEO PMB      |
| SIB F2S 4/4 | REV 05 | 710-022603 | DW7929 | F2S SIB      |
| B Board     | REV 05 | 710-023787 | DW7643 | NEO PMB      |

|             |        |            |        |                |
|-------------|--------|------------|--------|----------------|
| SIB F2S 4/6 | REV 05 | 710-022603 | DW7870 | F2S SIB        |
| B Board     | REV 05 | 710-023787 | DW7635 | NEO PMB        |
| Fan Tray 0  | REV 06 | 760-024497 | DV7831 | Front Fan Tray |
| Fan Tray 1  | REV 06 | 760-024497 | DV9614 | Front Fan Tray |
| Fan Tray 2  | REV 06 | 760-024502 | DV9618 | Rear Fan Tray  |
| Fan Tray 3  | REV 06 | 760-024502 | DV9616 | Rear Fan Tray  |
| Fan Tray 4  | REV 06 | 760-024502 | DV7807 | Rear Fan Tray  |
| Fan Tray 5  | REV 06 | 760-024502 | DV7828 | Rear Fan Tray  |

## Release Information

Command introduced before Junos OS Release 7.4.

**models** option introduced in Junos OS Release 8.2.

**sfc** option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

Information for **disk** and **usb** introduced in Junos OS Release 15.1X53-D60 for QFX10002, QFX10008, and QFX10016 switches.

**NOTE:** Routers and routing platforms use the basic syntax, unless otherwise listed. For example, the EX Series has an additional satellite parameter available.

## RELATED DOCUMENTATION

| [show chassis power](#)

## show chassis pic

### IN THIS SECTION

- [Syntax | 1275](#)
- [Syntax \(TX Matrix and TX Matrix Plus Routers\) | 1275](#)
- [Syntax \(MX Series Routers and EX Series Switches\) | 1275](#)

- [Syntax \(PTX Series Packet Transport Router and MX240, MX480, MX960, MX2010, and MX2020 Routers\) | 1276](#)
- [Syntax \(QFX Series\) | 1276](#)
- [Syntax \(ACX5048 and ACX5096 Routers\) | 1276](#)
- [Syntax \(ACX500 Routers\) | 1276](#)
- [Description | 1276](#)
- [Options | 1276](#)
- [Required Privilege Level | 1279](#)
- [Output Fields | 1279](#)
- [Sample Output | 1283](#)
- [Release Information | 1311](#)

## Syntax

```
show chassis pic fpc-slot slot-number pic-slot slot-number
```

## Syntax (TX Matrix and TX Matrix Plus Routers)

```
show chassis pic fpc-slot slot-number pic-slot slot-number
<lcc number>
```

## Syntax (MX Series Routers and EX Series Switches)

```
show chassis pic fpc-slot slot-number pic-slot slot-number
<all-members>
<local>
<member member-id>
```

## Syntax (PTX Series Packet Transport Router and MX240, MX480, MX960, MX2010, and MX2020 Routers)

```
show chassis pic transport fpc-slot slot-number pic-slot slot-number
```

## Syntax (QFX Series)

```
show chassis pic fpc-slot slot-number pic-slot slot-number
<interconnect-device name (fpc-slot slot-number | pic-slot slot-number)>
<node-device name pic-slot slot-number>
```

## Syntax (ACX5048 and ACX5096 Routers)

```
show chassis pic
(fpc-slot slot-number | pic-slot slot-number)
```

## Syntax (ACX500 Routers)

```
show chassis pic
(fpc-slot slot-number | pic-slot slot-number)
```

## Description

Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.

## Options

**fpc-slot *slot-number***

Display information about the PIC in this particular FPC slot:

- On a TX Matrix router, if you specify the number of the T640 router by using the **lcc *number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the `lcc number` option (the recommended method), replace *slot-number* with a value from 0 through 7. Otherwise, replace *slot-number* with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis pic fpc-slot 1 lcc 1 pic-slot 1
user@host> show chassis pic fpc-slot 9 pic-slot 1
```

- M120 routers only—Replace *slot-number* with a value from 0 through 5.
- MX80 routers only—Replace *slot-number* with a value from 0 through 1.
- MX104 routers only—Replace *slot-number* with a value from 0 through 2.
- MX240 routers only—Replace *slot-number* with a value from 0 through 2.
- MX480 routers only—Replace *slot-number* with a value from 0 through 5.
- MX960 routers only—Replace *slot-number* with a value from 0 through 11.
- MX2010 routers only—Replace *slot-number* with a value from 0 through 9.
- MX2020 routers only—Replace *slot-number* with a value from 0 through 19.
- MX2008 routers only—Replace *slot-number* with a value from 0 through 9.
- MX10003 routers only—Replace *slot-number* with a value from 0 through 1.
- Other routers—Replace *slot-number* with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace *slot-number* with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace *slot-number* with a value from 0 through 9 (switch's member ID).
  - EX8208 switches—Replace *slot-number* with a value from 0 through 7 (line card).
  - EX8216 switches—Replace *slot-number* with a value from 0 through 15 (line card).

- QFX Series:
  - QFX3500, QFX3600, QFX5100, and OCX Series standalone switches—Replace *slot-number* with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
  - QFabric systems—Replace *slot-number* with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

|                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>all-members</b>                     | (MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>interconnect-device <i>name</i></b> | (QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>lcc <i>number</i></b>               | <p>(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> <li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li> <li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li> </ul> |
| <b>local</b>                           | (MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>member <i>member-id</i></b>         | (MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value of 0 or 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>node-device <i>name</i></b>         | (QFabric systems only) (Optional) Display PIC information for a specified Node device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>pic-slot <i>slot-number</i></b> | Display information about the PIC in this particular PIC slot. For routers, replace <i>slot-number</i> with a value from 0 through 3. For EX3200 and EX4200 switches, replace <i>slot-number</i> with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace <i>slot-number</i> with 0. For the QFX3500 standalone switch and the QFabric system, replace <i>slot-number</i> with 0 or 1. |
| <b>transport</b>                   | Display PIC information for optical transport network.                                                                                                                                                                                                                                                                                                                                                                                               |

## Required Privilege Level

view

## Output Fields

Table 138 on page 1279 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

**Table 138: show chassis pic Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                      |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>                    | PIC type.<br><br><b>NOTE:</b> On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as <b>MIC-3D-1STM64-XFP</b> and with the SONET framing mode, the type is displayed as <b>MIC-3D-1OC192-XFP</b> . By default, the 1-port OC192/STM64 MICs displays the type as <b>MIC-3D-1OC192-XFP</b> . |
| <b>Account Layer2 Overhead</b> | (MX Series routers) Indicates whether functionality to count the Layer 2 overhead bytes in the interface statistics at the PIC level is enabled or disabled.                                                                                                                                                           |
| <b>ASIC type</b>               | Type of ASIC on the PIC.                                                                                                                                                                                                                                                                                               |

Table 138: show chassis pic Output Fields (Continued)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>State</b>       | <p>Status of the PIC. State is displayed only when a PIC is in the slot.</p> <ul style="list-style-type: none"> <li>• <b>Online</b>— PIC is online and running.</li> <li>• <b>Offline</b>—PIC is powered down.</li> <li>• <b>Empty</b>—No PIC is present.</li> <li>• <b>Present</b>—PIC is plugged in. The PIC is not powered on or operational.</li> <li>• <b>Onlining</b>—PIC is in the process of going online. PICs and rest of the hardware is initializing.</li> <li>• <b>Offlining</b>—PIC is in the process of going offline. PIC and rest of the hardware is being shutdown down to take the offline gracefully.</li> <li>• <b>Fault</b>—PIC is in an alarmed state and the PIC is not operational.</li> </ul> |
| <b>PIC version</b> | PIC hardware version.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Uptime</b>      | How long the PIC has been online.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Package</b>     | <p>(Multiservices interfaces cards only)</p> <p>The following services package are supported:</p> <ul style="list-style-type: none"> <li>• <b>Layer-2</b> (not applicable to the <a href="#">MS-MPC</a> and <a href="#">MS-MIC</a>.)</li> <li>• <b>Layer-3</b> (not applicable to the <a href="#">MS-MPC</a> and <a href="#">MS-MIC</a>.)</li> <li>• <b>extension-provider</b> (PICs of <a href="#">MS-MPC</a> and <a href="#">MS-MIC</a> support only this package.)</li> </ul>                                                                                                                                                                                                                                        |
| <b>Port Number</b> | Port number for the PIC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Cable Type</b>  | Type of cable connected to the port: <b>LH</b> , <b>LX</b> , or <b>SX</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 138: show chassis pic Output Fields (*Continued*)

| Field Name                                                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PIC Port Information<br/>(MX480 Router 100-Gigabit Ethernet CFP)</b> | Port-level information for the PIC. <ul style="list-style-type: none"> <li>• Port—Port number</li> <li>• Cable type—Type of optical transceiver installed.</li> <li>• Fiber type—Type of fiber. SM is single-mode.</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.</li> <li>• Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction</li> <li>• Xcvr Firmware—Transceiver firmware version.</li> </ul>                                                                                                                                                                                                             |
| <b>PIC Port Information<br/>(MX960 Router Bidirectional Optics )</b>    | Port-level information for the PIC. <ul style="list-style-type: none"> <li>• Port—Port number</li> <li>• Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D.</li> <li>• Fiber type—Type of fiber. SM is single-mode.</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.               <ul style="list-style-type: none"> <li>• BX10-10-km bidirectional optics.</li> <li>• BX40-40-km bidirectional optics.</li> <li>• SFP-LX-40-km SFP optics.</li> </ul> </li> <li>• Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.</li> </ul> |

Table 138: show chassis pic Output Fields (*Continued*)

| Field Name                                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PIC Port Information (Next-Generation SONET/SDH SFP)</b> | Port-level information for the next-generation SONET/SDH SFP PIC. <ul style="list-style-type: none"> <li>• Port—Port number.</li> <li>• Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed.</li> <li>• Fiber type—Type of fiber: <b>SM</b> (single-mode) or <b>MM</b> (multimode).</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.</li> <li>• Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.</li> </ul> |
| <b>PIC port information (MX104 router)</b>                  | Port-level information for the PIC. <ul style="list-style-type: none"> <li>• Port—Port number</li> <li>• Cable type—Type of optical transceiver installed.</li> <li>• Fiber type—Type of fiber. SM is single-mode.</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.</li> <li>• Wavelength—Wavelength of the transmitted signal.</li> <li>• Xcvr Firmware—Firmware version of the transceiver.</li> </ul>                                                                               |
| Port speed information                                      | Information pertaining to port speed: <ul style="list-style-type: none"> <li>• Port—Port number.</li> <li>• PFE—Packet Forwarding Engine slot number.</li> <li>• Capable Port Speed—Speed supported by each port.</li> </ul>                                                                                                                                                                                                                                                                                                                             |

Table 138: show chassis pic Output Fields (Continued)

| Field Name                  | Field Description                                                                                            |
|-----------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Multirate Mode</b>       | Rate-selectability status for the MIC: <b>Enabled</b> or <b>Disabled</b> .                                   |
| <b>Channelization</b>       | Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.                                   |
| <b>Administrative State</b> | Indicates the administrative state of the PIC. Possible values are: In Service (Default) and Out of Service. |
| <b>Operational State</b>    | Indicates the operational state of the PIC. Possible values are: Normal and Fault.                           |

## Sample Output

### show chassis pic fpc-slot pic-slot

```

user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
 Type 10x 1GE(LAN), 1000 BASE
 ASIC type H chip
 State Online
 PIC version 1.1
 Uptime 1 day, 50 minutes, 58 seconds
PIC Port Information:
 Port Cable Xcvr Xcvr Vendor
 Number Type Vendor Name Part Number
 0 GIGE 1000EX FINISAR CORP. FTRJ8519P1BNL-J3
 1 GIGE 1000EX FINISAR CORP. FTRJ-8519-7D-JUN

```

**show chassis pic fpc-slot pic-slot (PIC Offline)**

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
 State Offline

```

**show chassis pic fpc-slot pic-slot (FPC Offline)**

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online

```

**show chassis pic fpc-slot pic-slot (FPC Not Present)**

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty

```

**show chassis pic fpc-slot pic-slot (PIC Not Present)**

```

user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty

```

**show chassis pic fpc-slot 3 pic-slot 0 (M120 Router)**

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
 Type 2x G/E IQ, 1000 BASE
 ASIC type IQ GE 2 VLAN-TAG FPGA
 State Online
 PIC version 1.16
 Uptime 3 hours, 3 minutes

PIC Port Information:
 Port Cable Xcvr Xcvr Vendor
 Number Type Vendor Name Part Number

```

|   |             |               |                  |
|---|-------------|---------------|------------------|
| 0 | GIGE 1000SX | FINISAR CORP. | FTRJ8519P1BNL-J3 |
| 1 | GIGE 1000SX | FINISAR CORP. | FTRJ-8519-7D-JUN |

### show chassis pic fpc-slot pic-slot (MX150)

```

user@host> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
 Type Virtual
 State Online
 PIC version 0.0
 Uptime 7 days, 19 hours, 44 minutes, 40 seconds

PIC port information:
 Port Cable type Fiber Xcvr vendor Wave- Xcvr
 Cable type type Xcvr vendor part number length
 Firmware
 10 GIGE 1000T n/a Methode Elec. SP7041-M1-JN n/a
 0.0
 11 GIGE 1000T n/a Methode Elec. SP7041-M1-JN n/a
 0.0

```

### show chassis pic fpc-slot pic-slot (MX960 Router with Bidirectional Optics)

```

user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
 Type 10x 1GE (LAN)
 Account Layer2 Overhead Enabled
 State Online
 PIC version 0.0
 Uptime 18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:
 Port Cable type Fiber Xcvr vendor Wavelength
 Cable type type Xcvr vendor part number
 0 SFP-1000BASE-BX10-D SM SumitomoElectric SBP6H44-J3-BW-49 1490 nm
 1 SFP-1000BASE-BX10-D SM SumitomoElectric SBP6H44-J3-BW-49 1490 nm
 2 SFP-1000BASE-BX10-D SM SumitomoElectric SBP6H44-J3-BW-49 1490 nm
 3 SFP-1000BASE-BX10-D SM OCP TRXBG1LXDBVM2-JW 1490 nm
 4 SFP-1000BASE-BX10-D SM OCP TRXBG1LXDBVM2-JW 1490 nm

```

|   |                     |    |                  |                  |         |
|---|---------------------|----|------------------|------------------|---------|
| 5 | SFP-1000BASE-BX10-U | SM | SumitomoElectric | SBP6H44-J3-BW-31 | 1310 nm |
| 6 | SFP-1000BASE-BX10-U | SM | SumitomoElectric | SBP6H44-J3-BW-31 | 1310 nm |
| 7 | SFP-1000BASE-BX10-U | SM | OCP              | TRXBG1LXDDBMH-J1 | 1310 nm |
| 8 | SFP-1000BASE-BX10-U | SM | OCP              | TRXBG1LXDDBMH-J1 | 1310 nm |
| 9 | SFP-1000BASE-BX10-U | SM | SumitomoElectric | SBP6H44-J3-BW-31 | 1310 nm |

### show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
 Type 1X100GE CFP
 State Online
 PIC version 2.10
 Uptime 4 minutes, 48 seconds

PIC port information:
 Fiber Xcvr vendor
 Port Cable type type Xcvr vendor part number Wavelength
 0 100GBASE LR4 SM FINISAR CORP. FTLC1181RDNS-J3 1310 nm

 Xcvr vendor
 firmware version
 1.8

```

### show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
 Type AS-MXC
 State Online
 PIC version 1.0
 Uptime 11 hours, 18 minutes, 3 seconds

```

### show chassis pic fpc-slot pic-slot (MX960 Router with MPC5EQ)

```

user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:

```

```

Type 1X100GE CFP2 OTN
State Online
PIC version 0.0
Uptime 1 hour, 22 minutes, 42 seconds

PIC port information:

Port Cable type Fiber Xcvr vendor Wave- Xcvr
 type Xcvr vendor part number length
Firmware
0 100GBASE LR4 n/a Oclaro Inc. TRB5E20FNF-LF150 1309 nm
1.0

```

### show chassis pic fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
Type 1X100GE DWDM CFP2-ACO
State Online
PIC version 1.3
Uptime 9 hours, 4 minutes, 43 seconds

PIC port information:

Port Cable type Fiber Xcvr vendor Wave- Xcvr
 type Xcvr vendor part number length
Firmware
0 100G LH SM OCLARO TRB100AJ-01 1528.77 nm -
1568.36 nm 20.10

```

### show chassis pic fpc-slot pic-slot (MX10003 Routers)

```

user@host > show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 1 information:
Type MIC1
State Online
PIC version 1.5
Uptime 13 hours, 54 minutes, 33 seconds

```

## PIC port information:

| Port     | Cable type  | Fiber type | Xcvr vendor | Xcvr vendor part number | Wave-length | Xcvr |
|----------|-------------|------------|-------------|-------------------------|-------------|------|
| Firmware |             |            |             |                         |             |      |
| 0        | 40GBASE SR4 | MM         | AVAGO       | AFBR-79EQDZ-JU2         | 850 nm      | 0.0  |
| 11       | 40GBASE SR4 | MM         | AVAGO       | AFBR-79EQDZ-JU2         | 850 nm      | 0.0  |

## Port speed information:

| Port | PFE | Capable Port Speeds |
|------|-----|---------------------|
| 0    | 0   | 4x10GE, 40GE, 100GE |
| 1    | 0   | 4x10GE, 40GE, 100GE |
| 2    | 0   | 4x10GE, 40GE, 100GE |
| 3    | 0   | 4x10GE, 40GE, 100GE |
| 4    | 1   | 4x10GE, 40GE, 100GE |
| 5    | 1   | 4x10GE, 40GE, 100GE |
| 6    | 1   | 4x10GE, 40GE, 100GE |
| 7    | 1   | 4x10GE, 40GE, 100GE |
| 8    | 2   | 4x10GE, 40GE, 100GE |
| 9    | 2   | 4x10GE, 40GE, 100GE |
| 10   | 2   | 4x10GE, 40GE, 100GE |
| 11   | 2   | 4x10GE, 40GE, 100GE |

**show chassis pic fpc-slot pic-slot (PTX1000 and PTX10000)**

```
user@host > show chassis pic fpc-slot 0 pic-slot 0
```

```
FPC slot 0, PIC slot 0 information:
```

|             |                                        |
|-------------|----------------------------------------|
| Type        | 288X10GE/72X40GE/24X100GE              |
| State       | Online                                 |
| PIC version | 1.18                                   |
| Uptime      | 9 day, 5 hours, 10 minutes, 56 seconds |

## PIC port information:

| Port     | Cable type   | Fiber type | Xcvr vendor    | Xcvr vendor part number | Wave-length | Xcvr |
|----------|--------------|------------|----------------|-------------------------|-------------|------|
| Firmware |              |            |                |                         |             |      |
| 13       | 100GBASE LR4 | SM         | JUNIPER-SOURCE | SPQCELRCDFBJ2           | 1302 nm     | 0.0  |
| 25       | 100GBASE LR4 | SM         | JUNIPER-SOURCE | SPQCELRCDFAJ2           | 1302 nm     | 0.0  |
| 36       | 40GBASE LR4  | SM         | FINISAR CORP.  | FTL4C1QE1C-J1           | 1301 nm     |      |

```

0.0
 37 40GBASE LR4 SM FINISAR CORP. FTL4C1QE1C-J1 1301 nm
0.0
 54 40GBASE SR4 MM AVAGO AFBR-79EQDZ-JU1 850 nm
0.0

```

Port speed information:

| Port | PFE | Capable Port Speeds |
|------|-----|---------------------|
| 0    |     | 4x10GE, 40GE        |
| 1    |     | 4x10GE, 40GE, 100GE |
| 2    |     | 4x10GE, 40GE        |
| 3    |     | 4x10GE, 40GE        |
| 4    |     | 4x10GE, 40GE        |
| 5    |     | 4x10GE, 40GE, 100GE |
| 6    |     | 4x10GE, 40GE        |
| 7    |     | 4x10GE, 40GE, 100GE |
| 8    |     | 4x10GE, 40GE        |
| 9    |     | 4x10GE, 40GE        |
| 10   |     | 4x10GE, 40GE        |
| 11   |     | 4x10GE, 40GE, 100GE |
| 12   |     | 4x10GE, 40GE        |
| 13   |     | 4x10GE, 40GE, 100GE |
| 14   |     | 4x10GE, 40GE        |
| 15   |     | 4x10GE, 40GE        |
| 16   |     | 4x10GE, 40GE        |
| 17   |     | 4x10GE, 40GE, 100GE |
| 18   |     | 4x10GE, 40GE        |
| 19   |     | 4x10GE, 40GE, 100GE |
| 20   |     | 4x10GE, 40GE        |
| 21   |     | 4x10GE, 40GE        |
| 22   |     | 4x10GE, 40GE        |
| 23   |     | 4x10GE, 40GE, 100GE |
| 24   |     | 4x10GE, 40GE        |
| 25   |     | 4x10GE, 40GE, 100GE |
| 26   |     | 4x10GE, 40GE        |
| 27   |     | 4x10GE, 40GE        |
| 28   |     | 4x10GE, 40GE        |
| 29   |     | 4x10GE, 40GE, 100GE |
| 30   |     | 4x10GE, 40GE        |
| 31   |     | 4x10GE, 40GE, 100GE |
| 32   |     | 4x10GE, 40GE        |
| 33   |     | 4x10GE, 40GE        |
| 34   |     | 4x10GE, 40GE        |

```
35 4x10GE, 40GE, 100GE
36 4x10GE, 40GE
37 4x10GE, 40GE, 100GE
38 4x10GE, 40GE
39 4x10GE, 40GE
40 4x10GE, 40GE
41 4x10GE, 40GE, 100GE
42 4x10GE, 40GE
43 4x10GE, 40GE, 100GE
44 4x10GE, 40GE
45 4x10GE, 40GE
46 4x10GE, 40GE
47 4x10GE, 40GE, 100GE
48 4x10GE, 40GE
49 4x10GE, 40GE, 100GE
50 4x10GE, 40GE
51 4x10GE, 40GE
52 4x10GE, 40GE
53 4x10GE, 40GE, 100GE
54 4x10GE, 40GE
55 4x10GE, 40GE, 100GE
56 4x10GE, 40GE
57 4x10GE, 40GE
58 4x10GE, 40GE
59 4x10GE, 40GE, 100GE
60 4x10GE, 40GE
61 4x10GE, 40GE, 100GE
62 4x10GE, 40GE
63 4x10GE, 40GE
64 4x10GE, 40GE
65 4x10GE, 40GE, 100GE
66 4x10GE, 40GE
67 4x10GE, 40GE, 100GE
68 4x10GE, 40GE
69 4x10GE, 40GE
70 4x10GE, 40GE
```

### show chassis pic fpc-slot pic-slot (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

```
user@host > show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
```

```

Type 5X100GE DWDM CFP2-ACO
State Online
PIC version 1.17
Uptime 1 day, 5 hours, 15 minutes, 17 seconds

```

PIC port information:

| Port            | Cable type | Fiber type | Xcvr vendor     | Xcvr vendor part number | Wave-length | Xcvr |
|-----------------|------------|------------|-----------------|-------------------------|-------------|------|
| Firmware        |            |            |                 |                         |             |      |
| 0               | 100G LH    | SM         | MULTILANE SAL   | ML4030-ACO-2            | 1528.77 nm  | -    |
| 1568.36 nm 1.0  |            |            |                 |                         |             |      |
| 1               | 100G LH    | SM         | MULTILANE SAL   | ML4030-ACO-2            | 1528.77 nm  | -    |
| 1568.36 nm 1.0  |            |            |                 |                         |             |      |
| 2               | 100G LH    | SM         | JUNIPER-FUJITSU | FIM38500/222            | 1528.77 nm  | -    |
| 1568.36 nm 1.16 |            |            |                 |                         |             |      |
| 3               | 100G LH    | SM         | FUJITSU         | FIM38500/222            | 1528.77 nm  | -    |
| 1568.36 nm 1.16 |            |            |                 |                         |             |      |
| 4               | 100G LH    | SM         | FUJITSU         | FIM38500/222            | 1528.77 nm  | -    |
| 1568.36 nm 1.16 |            |            |                 |                         |             |      |

### show chassis pic fpc-slot pic-slot (MX480 Router with MPC4E)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
```

FPC slot 3, PIC slot 0 information:

```

Type 4x10GE SFPP
State Online
PIC version 0.0
Uptime 41 seconds

```

PIC port information:

| Port     | Cable type | Fiber type | Xcvr vendor  | Xcvr vendor part number | Wave-length | Xcvr |
|----------|------------|------------|--------------|-------------------------|-------------|------|
| Firmware |            |            |              |                         |             |      |
| 0        | 10GBASE SR | MM         | OPNEXT, INC. | TRS2001EM-0014          | 850 nm      |      |
| 0.0      |            |            |              |                         |             |      |
| 1        | 10GBASE SR | MM         | OPNEXT, INC. | TRS2001EM-0014          | 850 nm      |      |
| 0.0      |            |            |              |                         |             |      |

**show chassis pic fpc-slot pic-slot (MX480 router with OTN Interface)**

```

user@host> show chassis pci fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
 Type 12X10GE SFPP OTN
 State Online
 PIC version 0.0
 Uptime 5 hours, 28 minutes, 23 seconds

PIC port information:
 Port Cable type Fiber Xcvr vendor Wave- Xcvr
 type type Xcvr vendor part number length
Firmware
 0 10GBASE SR MM FINISAR CORP. FTLX8571D3BNL-J1 850 nm
0.0
 1 10GBASE SR MM FINISAR CORP. FTLX8571D3BCL-J1 850 nm
0.0
 2 10GBASE SR MM OPNEXT, INC. TRS2001EM-0014 850 nm
0.0

```

**show chassis pic fpc-slot pic-slot (MX2010 Router with OTN Interfaces)**

```

user@host> show chassis pic fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
 Type 2X100GE CFP2 OTN
 State Online
 PIC version 1.9
 Uptime 3 hours, 56 minutes, 16 seconds

PIC port information:
 Port Cable type Fiber Xcvr vendor Wave- Xcvr
 type type Xcvr vendor part number length
Firmware
 0 100GBASE LR4-D SM FUJITSU FIM37300/222 1310 nm
1.3
 1 100GBASE SR10 MM AVAGO AFBR-8420Z n/a 1.0

```

**show chassis pic fpc-slot pic-slot (MX2010 Router)**

```

user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
 Type 1X100GE CFP
 Account Layer2 Overhead Enabled
 State Online
 PIC version 0.0
 Uptime 14 hours, 51 seconds

```

**show chassis pic fpc-slot pic-slot (MX2020 Router)**

```

user@host> show chassis pic fpc-slot 19 pic-slot 3
FPC slot 19, PIC slot 3 information:
 Type 4x 10GE(LAN) SFP+
 Account Layer2 Overhead Enabled
 State Online
 PIC version 0.0
 Uptime 1 day, 11 hours, 26 minutes, 36 seconds

PIC port information:

```

| Port | Cable type | Fiber type | Xcvr vendor      | Xcvr vendor part number | Wave-length | Xcvr |
|------|------------|------------|------------------|-------------------------|-------------|------|
| 0    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M          | 850 nm      |      |
| 1    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M          | 850 nm      |      |
| 2    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M          | 850 nm      |      |
| 3    | 10GBASE SR | MM         | SumitomoElectric | SPP5200SR-J6-M          | 850 nm      |      |

```

Firmware
0.0
0.0
0.0
0.0

```

**show chassis pic fpc-slot pic-slot (MX2020 Router with MPC5EQ and MPC6E)**

```

user@host> show chassis pic fpc-slot 18 pic-slot 2
FPC slot 18, PIC slot 2 information:
 Type 3X40GE QSFP
 State Online

```

```

PIC version 0.0
Uptime 6 minutes, 31 seconds

PIC port information:

Port Cable type Fiber Xcvr vendor Wave- Xcvr
 type Xcvr vendor part number length
Firmware
0 40GBASE SR4 MM AVAGO AFBR-79E4Z-D-JU2 850 nm
0.0
1 40GBASE SR4 MM AVAGO AFBR-79E4Z-D-JU2 850 nm
0.0
2 40GBASE SR4 MM AVAGO AFBR-79E4Z-D-JU2 850 nm
0.0

```

### show chassis pic fpc-slot pic-slot (MX2020 Router with MPC6E and OTN MIC)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 0, PIC slot 1 information:
Type 24X10GE SFPP OTN
State Online
PIC version 1.1
Uptime 1 hour, 33 minutes, 59 seconds

PIC port information:

Port Cable type Fiber Xcvr vendor Wave- Xcvr
 type Xcvr vendor part number length
Firmware
7 10GBASE SR MM SumitomoElectric SPP5200SR-J6-M 850 nm
0.0
9 10GBASE SR MM FINISAR CORP. FTLX8571D3BNL-J1 850 nm
0.0
12 10GBASE LR SM FINISAR CORP. FTLX1472M3BNL-J3 1310 nm
0.0
20 10GBASE ZR SM FINISAR CORP. FTLX1871M3BNL-J3 1550 nm
0.0
21 10GBASE ER SM FINISAR CORP. FTLX1671D3BTL-J4 1550 nm
0.0
22 10GBASE LR SM SOURCEPHOTONICS SPP10SLREDFCJNP 1310 nm
0.0
23 10GBASE LR SM FINISAR CORP. FTLX1471D3BNL-J1 1310 nm 0.0

```

**show chassis pic fpc-slot pic-slot (MX2020 Router with MPC4E)**

```

user@host> show chassis pic fpc-slot 14 pic-slot 0
FPC slot 14, PIC slot 2 information:
 Type 4x10GE SFPP
 State Online
 PIC version 0.0
 Uptime 1 day, 14 hours, 49 minutes, 9 seconds

PIC port information:
 Port Cable type Fiber Xcvr vendor Wave- Xcvr
 type Xcvr vendor part number length
Firmware
 0 10GBASE SR MM SumitomoElectric SPP5100SR-J3 850 nm
0.0
 1 10GBASE SR MM SumitomoElectric SPP5100SR-J3 850 nm
0.0
 3 10GBASE SR MM SumitomoElectric SPP5100SR-J3 850 nm
0.0

```

**show chassis pic fpc-slot pic-slot (MX2010 Router)**

```

user@host> show chassis pic fpc-slot 9 pic-slot 3
FPC slot 9, PIC slot 3 information:
 Type 1X100GE CFP
 Account Layer2 Overhead Enabled
 State Online
 PIC version 0.0
 Uptime 14 hours, 51 seconds

```

**show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)**

```

user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
 Type 100GE SLOT1
 ASIC type Brooklyn 100GE FPGA
 State Online
 PIC version 1.3
 Uptime 10 minutes, 44 seconds

```

PIC port information:

| Port | Cable type  | Fiber type | Xcvr vendor | Xcvr vendor part number | Wavelength |
|------|-------------|------------|-------------|-------------------------|------------|
| 0    | 10GBASE LR4 | SM         | Opnext Inc. | TRC5E20ENFSF000F        | 1310 nm    |

### show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```
user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
```

lcc0-re0:

-----  
PIC fpc slot 1 pic slot 1 information:

|             |                                        |
|-------------|----------------------------------------|
| Type        | 4x OC-3 SONET, SMIR                    |
| ASIC type   | D chip                                 |
| State       | Online                                 |
| PIC version | 1.2                                    |
| Uptime      | 5 days, 2 hours, 12 minutes, 8 seconds |

### show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```
user@host> show chassis pic pic-slot 0 fpc-slot 8
```

lcc0-re0:

-----  
FPC slot 8, PIC slot 0 information:

|        |                                 |
|--------|---------------------------------|
| Type   | 1x 10GE (LAN/WAN)               |
| State  | Online                          |
| Uptime | 2 hours, 46 minutes, 23 seconds |

PIC port information:

| Port | Cable type | Fiber type | Xcvr vendor   | Xcvr vendor part number | Wavelength |
|------|------------|------------|---------------|-------------------------|------------|
| 0    | 10GBASE ZR | SM         | Opnext Inc.   | TRF7061BN-LF150         | 1550 nm    |
| 0    | 10GBASE ZR | SM         | FINISAR CORP. | FTRX-1811-3-J2          | 1550 nm    |

### show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```
user@host> show chassis pic fpc-slot 4 pic-slot 0
```

FPC slot 4, PIC slot 0 information:

```

Type 4x OC-3 1x OC-12 SFP
ASIC type D FPGA
State Online
PIC version 1.3
Uptime 1 day, 50 minutes, 4 seconds

```

PIC port information:

| Port | Cable type       | Fiber type | Xcvr vendor   | Xcvr vendor part number | Wavelength |
|------|------------------|------------|---------------|-------------------------|------------|
| 0    | OC48 short reach | SM         | FINISAR CORP. | FTRJ1321P1BTL-J2        | 1310 nm    |
| 1    | OC3 short reach  | MM         | OCP           | TRPA03MM3BAS-JE         | 1310 nm    |
| 2    | OC3 short reach  | MM         | OCP           | TRXA03MM3BAS-JW         | 1310 nm    |
| 3    | OC12 inter reach | SM         | FINISAR CORP. | FTLF1322P1BTR           | 1310 nm    |

### show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```

user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
Type 12x T1/E1 CE
State Online
PIC version 1.1
CPU load average 1 percent
Interrupt load average 0 percent
Total DRAM size 128 MB
Memory buffer utilization 100 percent
Memory heap utilization 4 percent
Uptime 1 day, 22 hours, 28 minutes, 12 seconds
Internal Clock Synchronization Normal

```

### show chassis pic fpc-slot 0 pic-slot 1 (4x CHOC3 SONET CE SFP)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
Type 4x CHOC3 SONET CE SFP
State Online
PIC version 1.3
CPU load average 1 percent
Interrupt load average 0 percent
Total DRAM size 128 MB

```

```

Memory buffer utilization 99 percent
Memory heap utilization 4 percent
Uptime 1 day, 22 hours, 55 minutes, 37 seconds
Internal Clock Synchronization Normal

```

PIC port information:

| Port | Cable type      | Fiber type | Xcvr vendor | Xcvr vendor part number | Wavelength |
|------|-----------------|------------|-------------|-------------------------|------------|
| 0    | OC3 short reach | MM         | AVAGO       | HFBR-57E0P-JU2          | n/a        |
| 1    | OC3 short reach | MM         | AVAGO       | HFBR-57E0P-JU2          | n/a        |
| 3    | OC3 long reach  | SM         | OPNEXT INC  | TRF5456AVLB314          | 1310 nm    |

### show chassis pic fpc-slot 0 pic-slot 0 (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
```

FPC slot 0, PIC slot 0 information:

```

Type MIC-3D-8OC3OC12-4OC48
State Online
PIC version 1.8
Uptime 3 days, 22 hours, 3 minutes, 50 seconds

```

PIC port information:

| Port | Cable type       | Fiber type | Xcvr vendor  | Xcvr vendor part number | Wavelength |
|------|------------------|------------|--------------|-------------------------|------------|
| 1    | OC12 inter reach | SM         | FINISAR CORP | FTRJ1322P1BTR-J3        | 1310 nm    |
| 7    | OC12 inter reach | SM         | FINISAR CORP | FTRJ1322P1BTR-J3        | 1310 nm    |

Multirate Mode Enabled

### show chassis pic fpc-slot 3 pic-slot 0 (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 3 pic-slot 0
```

FPC slot 3, PIC slot 0 information:

```

Type MIC-3D-8CHOC3-4CHOC12
State Online
PIC version 1.9
Uptime 1 hour, 21 minutes, 24 seconds

```

PIC port information:

| Port | Cable type       | Fiber |               | Xcvr vendor      |  | Wavelength |
|------|------------------|-------|---------------|------------------|--|------------|
|      |                  | type  | Xcvr vendor   | part number      |  |            |
| 0    | OC12 short reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |
| 1    | OC12 short reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |
| 2    | OC12 inter reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J2 |  | 1310 nm    |
| 4    | OC12 short reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |
| 5    | OC12 short reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |
| 6    | OC12 short reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |
| 7    | OC12 short reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |

### show chassis pic fpc-slot 5 pic-slot 0 (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```

user@host> show chassis pic fpc-slot 5 pic-slot 0
FPC slot 5, PIC slot 0 information:
 Type MIC-3D-4CHOC3-2CHOC12
 State Online
 PIC version 1.9
 Uptime 1 hour, 21 minutes

PIC port information:

```

| Port | Cable type       | Fiber |               | Xcvr vendor      |  | Wavelength |
|------|------------------|-------|---------------|------------------|--|------------|
|      |                  | type  | Xcvr vendor   | part number      |  |            |
| 1    | OC12 inter reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |
| 2    | OC12 inter reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |
| 3    | OC12 short reach | SM    | FINISAR CORP. | FTRJ1322P1BTR-J3 |  | 1310 nm    |

### show chassis pic fpc-slot 1 pic-slot 0 (1-port OC192/STM64 MIC with XFP)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
 Type MIC-3D-1OC192-XFP
 State Online
 PIC version 1.2
 Uptime 1 day, 11 hours, 4 minutes, 6 seconds

PIC port information:

```

| Fiber |  | Xcvr vendor |  |
|-------|--|-------------|--|
|-------|--|-------------|--|

| Port | Cable type        | type | Xcvr vendor   | part number      | Wavelength |
|------|-------------------|------|---------------|------------------|------------|
| 0    | OC192 short reach | n/a  | FINISAR CORP. | FTLX1412M3BCL-J3 | 1310 nm    |

### show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
 Type MIC-3D-8DS3-E3
 State Online
 PIC version 1.10
 Uptime 4 days, 1 hour, 29 minutes, 19 seconds
 Channelization Mode Disabled

```

### show chassis pic fpc-slot pic-slot (OTN)

```

user@host> show chassis pic fpc-slot 5 pic-slot 0
PIC fpc slot 5 pic slot 0 information:
 Type 1x10GE(LAN),OTN
 ASIC type H chip
 State Online
 PIC version 1.0
 Uptime 5 minutes, 50 seconds

```

### show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```

user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
 Type 48x 10G-SFP+ Builtin
 State Online
 Uptime 3 days, 3 hours, 5 minutes, 20 seconds

```

### show chassis pic fpc-slot pic-slot (QFX5100 Switches and OCX Series )

```

user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
 Type Unknown Builtin

```

```

State Online
Uptime 1 day, 17 hours, 5 minutes, 9 seconds

```

### show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```

user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
Type 16x 40G-GEBuiltIn
State Online
Uptime 2 hours, 47 minutes, 40 seconds

```

### show chassis pic node-device fpc-slot pic-slot (QFabric System)

```

user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
Type 48x 10G-SFP+BuiltIn
State Online
Uptime 2 hours, 52 minutes, 37 seconds

PIC port information:

```

| Port | Cable type | Fiber type | Xcvr vendor      | Xcvr vendor part number | Wavelength |
|------|------------|------------|------------------|-------------------------|------------|
| 0    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 1    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 2    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 3    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 4    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 5    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 6    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 7    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 8    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 9    | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 10   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 11   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 12   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 13   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 14   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 15   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |
| 16   | 10GBASE SR | MM         | SumitomoElectric | SPP5101SR-J3            | 850 nm     |

|    |            |    |                  |              |        |
|----|------------|----|------------------|--------------|--------|
| 17 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 18 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 19 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 20 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 21 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 22 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 23 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 24 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 25 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 26 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 27 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 28 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 29 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 30 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 31 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 32 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 33 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 34 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 35 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 36 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 37 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 38 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 39 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 40 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 41 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 42 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 43 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 44 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 45 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 46 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |
| 47 | 10GBASE SR | MM | SumitomoElectric | SPP5101SR-J3 | 850 nm |

### show chassis pic fpc-slot 0 pic-slot 1 (ACX2000 Universal Access Router)

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
 Type 8x 1GE(LAN) RJ45 Builtin
 State Online
 Uptime 6 days, 2 hours, 51 minutes, 11 seconds

```

**show chassis pic FPC-slot 1 PIC-slot 0 (MX Routers with Media Services Blade [MSB])**

```

user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
 Type AS-MSB
 State Online
 PIC version 1.6
 Uptime 11 hours, 17 minutes, 56 seconds

```

**show chassis pic FPC slot 1, PIC slot 2 (MX Routers with Media Services Blade [MSB])**

```

user@switch> show chassis pic fpc-slot 1 pic-slot 2
Type AS-MXC
 State Online
 PIC version 1.0
 Uptime 11 hours, 18 minutes, 3 seconds

```

**show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)**

```

user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State: In Service
Operational State: Normal

```

**show chassis pic transport fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)**

```

user@host> show chassis pic transport fpc-slot 3 pic-slot 0
Administrative State: In Service
Operational State: Normal

```

**show chassis pic fpc-slot 0 pic-slot 0 (ACX5096 Router)**

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
```

```
FPC slot 0, PIC slot 0 information:
```

```

Type 96x10G-8x40G
State Online
PIC version 2.9
Uptime 21 hours, 28 minutes, 13 seconds

```

```
PIC port information:
```

| Port | Cable type | Fiber type | Xcvr vendor      | Xcvr vendor part number | Wave-length | Xcvr |
|------|------------|------------|------------------|-------------------------|-------------|------|
| 0    | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BCL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 1    | 10GBASE LR | SM         | FINISAR CORP.    | FTLX1471D3BCL-J1        | 1310 nm     |      |
| 0.0  |            |            |                  |                         |             |      |
| 3    | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 4    | 10GBASE SR | MM         | SumitomoElectric | SPP5100SR-J3            | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 5    | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 6    | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 7    | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 8    | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 9    | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 10   | 10GBASE SR | MM         | OPNEXT, INC.     | TRS2001EN-0014          | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 11   | 10GBASE SR | MM         | SumitomoElectric | SPP5100SR-J3            | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 12   | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 13   | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BCL-J1        | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 14   | 10GBASE SR | MM         | SumitomoElectric | SPP5100SR-J3            | 850 nm      |      |
| 0.0  |            |            |                  |                         |             |      |
| 15   | 10GBASE SR | MM         | FINISAR CORP.    | FTLX8571D3BNL-J1        | 850 nm      |      |

|     |             |    |                  |                  |         |  |
|-----|-------------|----|------------------|------------------|---------|--|
| 0.0 |             |    |                  |                  |         |  |
| 16  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 17  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 18  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 19  | 10GBASE LR  | SM | FINISAR CORP.    | FTLX1471D3BCL-J1 | 1310 nm |  |
| 0.0 |             |    |                  |                  |         |  |
| 20  | 10GBASE LR  | SM | FINISAR CORP.    | FTLX1471D3BNL-J1 | 1310 nm |  |
| 0.0 |             |    |                  |                  |         |  |
| 21  | 10GBASE SR  | MM | SumitomoElectric | SPP5100SR-J3     | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 22  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 23  | 10GBASE SR  | MM | SumitomoElectric | SPP5100SR-J3     | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 24  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 25  | 10GBASE USR | MM | FINISAR CORP.    | FTLX8570D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 26  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 27  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 28  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 29  | GIGE 1000SX | MM | FINISAR CORP.    | FTLF8519P3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 31  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 32  | 10GBASE SR  | MM | SumitomoElectric | SPP5100SR-J3     | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 33  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 34  | 10GBASE SR  | MM | SumitomoElectric | SPP5100SR-J3     | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 35  | 10GBASE USR | MM | FINISAR CORP.    | FTLX8570D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 36  | 10GBASE USR | MM | FINISAR CORP.    | FTLX8570D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 37  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |

|     |               |    |                  |                  |         |
|-----|---------------|----|------------------|------------------|---------|
| 38  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 40  | GIGE 1000LX10 | SM | FINISAR CORP.    | FTLF1318P2BTL-J1 | 1310 nm |
| 0.0 |               |    |                  |                  |         |
| 41  | 10GBASE LR    | SM | OPNEXT, INC      | TRS5021EN-S201   | 1310 nm |
| 0.0 |               |    |                  |                  |         |
| 42  | 10GBASE LR    | SM | FINISAR CORP.    | FTLX1471D3BCL-J1 | 1310 nm |
| 0.0 |               |    |                  |                  |         |
| 43  | 10GBASE LR    | SM | SumitomoElectric | SPP5100LR-J3     | 1310 nm |
| 0.0 |               |    |                  |                  |         |
| 44  | 10GBASE LR    | SM | SumitomoElectric | SPP5100LR-J3     | 1310 nm |
| 0.0 |               |    |                  |                  |         |
| 45  | 10GBASE LR    | SM | FINISAR CORP.    | FTLX1471D3BCL-J1 | 1310 nm |
| 0.0 |               |    |                  |                  |         |
| 46  | 10GBASE LR    | SM | FINISAR CORP.    | FTLX1471D3BCL-J1 | 1310 nm |
| 0.0 |               |    |                  |                  |         |
| 47  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 48  | 10GBASE SR    | MM | SumitomoElectric | SPP5100SR-J3     | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 49  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 50  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 51  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 52  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 53  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 54  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 55  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 56  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 57  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 58  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 59  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |
| 0.0 |               |    |                  |                  |         |
| 60  | 10GBASE SR    | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |

|     |             |    |                  |                  |         |  |
|-----|-------------|----|------------------|------------------|---------|--|
| 0.0 |             |    |                  |                  |         |  |
| 61  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 62  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 63  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 64  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 65  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 66  | 10GBASE SR  | MM | SumitomoElectric | SPP5200SR-J6-M   | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 67  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 68  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 69  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 70  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 71  | 10GBASE LR  | SM | FINISAR CORP.    | FTLX1471D3BNL-J1 | 1310 nm |  |
| 0.0 |             |    |                  |                  |         |  |
| 72  | 10GBASE LR  | SM | FINISAR CORP.    | FTLX1471D3BCL-J1 | 1310 nm |  |
| 0.0 |             |    |                  |                  |         |  |
| 73  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 74  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 75  | 10GBASE SR  | MM | FINISAR CORP.    | FTLX8571D3BCL-J1 | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 76  | 10GBASE SR  | MM | SumitomoElectric | SPP5100SR-J3     | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 77  | 10GBASE USR | MM | OPNEXT, INC.     | TRS20A0EN-0014   | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 78  | 10GBASE USR | MM | OPNEXT, INC.     | TRS20A0EN-0014   | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |
| 79  | 10GBASE LRM | MM | OPNEXT INC       | TRS5001EN-0014   | 1310 nm |  |
| 0.0 |             |    |                  |                  |         |  |
| 80  | 10GBASE LRM | MM | OPNEXT INC       | TRS5001EN-0014   | 1310 nm |  |
| 0.0 |             |    |                  |                  |         |  |
| 81  | 10GBASE USR | MM | OPNEXT, INC.     | TRS20A0EN-0014   | 850 nm  |  |
| 0.0 |             |    |                  |                  |         |  |

|     |         |       |     |                  |                  |         |
|-----|---------|-------|-----|------------------|------------------|---------|
| 82  | 10GBASE | USR   | MM  | OPNEXT, INC.     | TRS20A0EN-0014   | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 83  | 10GBASE | USR   | MM  | OPNEXT, INC.     | TRS20A0EN-0014   | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 84  | 10GBASE | USR   | MM  | OPNEXT, INC.     | TRS20A0EN-0014   | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 85  | 10GBASE | LR    | SM  | OPNEXT, INC      | TRS5021EN-S201   | 1310 nm |
| 0.0 |         |       |     |                  |                  |         |
| 86  | 10GBASE | ER    | SM  | OPNEXT, INC      | TRS7050EN-S201   | 1550 nm |
| 0.0 |         |       |     |                  |                  |         |
| 87  | 10GBASE | LRM   | MM  | OPNEXT INC       | TRS5001EN-0014   | 1310 nm |
| 0.0 |         |       |     |                  |                  |         |
| 88  | 10GBASE | LRM   | MM  | OPNEXT INC       | TRS5001EN-0014   | 1310 nm |
| 0.0 |         |       |     |                  |                  |         |
| 89  | 10GBASE | LRM   | MM  | OPNEXT INC       | TRS5001EN-0014   | 1310 nm |
| 0.0 |         |       |     |                  |                  |         |
| 90  | 10GBASE | LRM   | MM  | OPNEXT INC       | TRS5001EN-0014   | 1310 nm |
| 0.0 |         |       |     |                  |                  |         |
| 91  | 10GBASE | USR   | MM  | FINISAR CORP.    | FTLX8570D3BCL-J1 | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 92  | 10GBASE | USR   | MM  | FINISAR CORP.    | FTLX8570D3BCL-J1 | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 93  | 10GBASE | LR    | SM  | SumitomoElectric | SPP5100LR-J3     | 1310 nm |
| 0.0 |         |       |     |                  |                  |         |
| 94  | 10GBASE | LR    | SM  | FINISAR CORP.    | FTLX1471D3BNL-J1 | 1310 nm |
| 0.0 |         |       |     |                  |                  |         |
| 95  | 10GBASE | SR    | MM  | FINISAR CORP.    | FTLX8571D3BNL-J1 | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 96  | 40GBASE | SR4   | MM  | AVAGO            | AFBR-79E4Z-D-JU1 | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 97  | 40GBASE | SR4   | MM  | AVAGO            | AFBR-79E4Z-D-JU1 | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 98  | 40GBASE | SR4   | MM  | AVAGO            | AFBR-79EQDZ-JU1  | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 99  | 40GBASE | SR4   | MM  | AVAGO            | AFBR-79EQDZ-JU1  | 850 nm  |
| 0.0 |         |       |     |                  |                  |         |
| 100 | 40GBASE | CU 1M | n/a | Molex Inc.       | 1110409055       | n/a     |
| 0.0 |         |       |     |                  |                  |         |
| 101 | 40GBASE | CU 1M | n/a | Molex Inc.       | 1110409055       | n/a     |
| 0.0 |         |       |     |                  |                  |         |
| 102 | 40GBASE | CU 1M | n/a | Molex Inc.       | 1110409055       | n/a     |
| 0.0 |         |       |     |                  |                  |         |
| 103 | 40GBASE | CU 1M | n/a | Molex Inc.       | 1110409055       | n/a     |



**show chassis pic fpc-slot 0 pic-slot 1 (ACX500 Router)**

```

user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
 Type 4x 1GE(LAN) RJ45, SFP Built-in
 State Online
 Uptime 17 hours, 54 minutes, 45 seconds

```

**show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)**

```

user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State: In Service
Operational State: Normal

```

**show chassis pic transport fpc-slot pic-slot (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)**

```

user@host> show chassis pic transport fpc-slot 3 pic-slot 0
Administrative State: In Service
Operational State: Normal

```

**show chassis pic fpc-slot 7 pic-slot 1 (MX960 Router MPC10E-15C-MRATE Line Card)**

```

user@router> show chassis pic fpc-slot 7 pic-slot 1

FPC slot 7, PIC slot 1 information:
 Type MRATE-5xQSFP
 State Online
 PIC version 0.0
 Uptime 3 hours, 33 minutes, 21 seconds

PIC port information:

```

| Xcvr     | JNPR         | Fiber | Xcvr vendor     | Wave-           |         |
|----------|--------------|-------|-----------------|-----------------|---------|
| Port     | Cable type   | type  | Xcvr vendor     | part number     | length  |
| Firmware | Rev          |       |                 |                 |         |
| 0        | 100GBASE LR4 | SM    | JUNIPER-FINISAR | FTLC1151RDPL-J3 | 1302 nm |

0.0 REV 01

Port speed information:

| Port | PFE | Capable Port Speeds |
|------|-----|---------------------|
| 0    | 1   | 4x10GE, 40GE, 100GE |
| 1    | 1   | 4x10GE, 40GE, 100GE |
| 2    | 1   | 4x10GE, 40GE, 100GE |
| 3    | 1   | 4x10GE, 40GE, 100GE |
| 4    | 1   | 4x10GE, 40GE, 100GE |

## Release Information

Command introduced before Junos OS Release 7.4.

**transport** option introduced in Junos OS Release 16.1R1 for MX Series Routers.

## RELATED DOCUMENTATION

[request chassis pic](#)

[show chassis hardware | 1253](#)

[100-Gigabit Ethernet Type 4 PIC with CFP Overview | 224](#)

## show ethernet-switching redundancy-groups

### IN THIS SECTION

- [Syntax | 1312](#)
- [Description | 1312](#)
- [Options | 1312](#)
- [Required Privilege Level | 1312](#)
- [Output Fields | 1312](#)
- [Sample Output | 1315](#)
- [Release Information | 1316](#)

## Syntax

```
show ethernet-switching redundancy-groups
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs
```

## Description

Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or switch. Note that the Redundancy Group ID is inherited by the bridging domain or VLAN from member AE interfaces.

## Options

|                            |                                                                                                                                                                                                              |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>redundancy-group-id</i> | (Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group. |
| <i>arp-statistics</i>      | (Optional) Count of ARP packets sent and received by the two MC-AE nodes.                                                                                                                                    |
| <i>nd-statistics</i>       | (Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.                                                                                                                     |
| <i>remote-macs</i>         | (Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.                                                                                                     |

## Required Privilege Level

view

## Output Fields

Output fields are listed in the approximate order in which they appear.

**Table 139: show ethernet-switching redundancy-groups arp-statistics Output Fields**

| Field Name                        | Field Description                                                              |
|-----------------------------------|--------------------------------------------------------------------------------|
| Redundancy Group ID               | Redundancy Group to which the following details apply.                         |
| MCLAG ARP Statistics Group ID     | ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance. |
| ARP Rx Count From Line            | Total number of ARPs received from the Line.                                   |
| ARP Tx Count To Peer              | Total number of ARPs sent to the peer.                                         |
| ARP Rx Count From Peer            | Total number of ARPs received from the peer.                                   |
| ARP Drop Count received from line | Total number of ARPs sent by the peer that were received.                      |
| ARP Drop Count received from peer | Total number of ARPs sent by the peer that were dropped                        |
| ARP Install Count                 | ARP Install Count                                                              |

**Table 140: show ethernet-switching redundancy-groups nd-statistics Output Fields**

| Field Name          | Field Description                                      |
|---------------------|--------------------------------------------------------|
| Redundancy Group ID | Redundancy Group to which the following details apply. |

**Table 140: show ethernet-switching redundancy-groups nd-statistics Output Fields (Continued)**

| Field Name                       | Field Description                                                                            |
|----------------------------------|----------------------------------------------------------------------------------------------|
| MCLAG ND Statistics Group ID     | Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MCLAG) instance. |
| ND Rx Count From Line            | Total number of Neighbor Discovery packets received from the Line.                           |
| ND Tx Count To Peer              | Total number of Neighbor Discovery packets sent to the peer.                                 |
| ND Rx Count From Peer            | Total number of Neighbor Discovery packets received from the peer.                           |
| ND Drop Count received from line | Total number of Neighbor Discovery packets sent by the peer that were received.              |
| ND Drop Count received from peer | Total number of Neighbor Discovery packets sent by the peer that were dropped                |
| ND Install Count                 | ND Install Count                                                                             |

**Table 141: show ethernet-switching redundancy-groups remote-macs Output Fields**

| Field Name          | Field Description                                      |
|---------------------|--------------------------------------------------------|
| Redundancy Group ID | Redundancy Group to which the following details apply. |
| Service ID          | Service ID (configured at the routing instance level). |
| Peer-Addr           | IP address of the remote peer.                         |

**Table 141: show ethernet-switching redundancy-groups remote-macs Output Fields (Continued)**

| Field Name | Field Description                                                                                                  |
|------------|--------------------------------------------------------------------------------------------------------------------|
| VLAN       | Virtual LAN identifier associated with the redundancy group.                                                       |
| MAC        | Hardware media access control address associated with the redundancy group.                                        |
| MCAE-ID    | ID number of the MC-AE used by the redundancy group.                                                               |
| Flags      | Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected. |
| Status     | Installation state: Installed or Not Installed.                                                                    |

## Sample Output

### show ethernet-switching redundancy-groups arp-statistics

```
user@host> show ethernet-switching redundancy-groups arp-statistics
```

```
Redundancy Group ID : 1 Flags : Local Connect,Remote Connect
```

#### MCLAG ARP Statistics

```
Group ID : 1
ARP Rx Count From Line : 3493
ARP Tx Count To Peer : 647
ARP Rx Count From Peer : 0
ARP Install Count : 0
ARP Drop Count received from line : 2846
ARP Drop Count received from peer : 0
```

## show ethernet-switching redundancy-groups nd-statistics

```

user@host> show ethernet-switching redundancy-groups nd-statistics

Redundancy Group ID : 1 Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID : 1
ND Rx Count From Line : 52
ND Tx Count To Peer : 15
ND Rx Count From Peer : 39
ND Install Count : 34
ND Drop Count received from line : 37
ND Drop Count received from peer : 5

```

## show ethernet-switching redundancy-groups remote-macs

```

user@host> show ethernet-switching redundancy-groups <redundancy-group-id> remote-macs

Redundancy Group ID : 1 Flags : Local Connect,Remote Connect

Service-id Peer-Addr VLAN MAC MCAE-ID Subunit
Opcode Flags Status
10 0 10.3.3.2 100 80:ac:ac:1f:10:a1 1 0
1 0 Installed

```

## show ethernet-switching redundancy-groups group-id

```

user@host> show ethernet-switching redundancy-groups 1

Redundancy Group ID : 1 Flags : Local Connect,Remote Connect

```

## Release Information

Command introduced in Junos OS Release 13.2.

## RELATED DOCUMENTATION

[Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up](#)

## show interfaces (Adaptive Services)

### IN THIS SECTION

- [Syntax | 1317](#)
- [Description | 1317](#)
- [Options | 1317](#)
- [Required Privilege Level | 1318](#)
- [Output Fields | 1318](#)
- [Sample Output | 1324](#)
- [Release Information | 1327](#)

### Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Description

Display status information about the specified adaptive services interface.

### Options

*interface-type*

On M Series and T Series routers, the interface type is **sp- *fpc/port***.

|                                           |                                                                               |
|-------------------------------------------|-------------------------------------------------------------------------------|
| <b>brief   detail   extensive   terse</b> | (Optional) Display the specified level of output.                             |
| <b>descriptions</b>                       | (Optional) Display interface description strings.                             |
| <b>media</b>                              | (Optional) Display media-specific information about network interfaces.       |
| <b>snmp-index <i>snmp-index</i></b>       | (Optional) Display information for the specified SNMP index of the interface. |
| <b>statistics</b>                         | (Optional) Display static interface statistics.                               |

## Required Privilege Level

view

## Output Fields

[Table 142 on page 1318](#) lists the output fields for the **show interfaces** (adaptive services and redundant adaptive services) command. Output fields are listed in the approximate order in which they appear.

**Table 142: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields**

| Field Name                | Field Description                                                                                                                 |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical Interface</b> |                                                                                                                                   |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                   |
| <b>Enabled</b>            | State of the interface. Possible values are described in the “Enabled Field” section of <i>Common Output Fields Description</i> . |
| <b>Interface index</b>    | Physical interface's index number, which reflects its initialization sequence.                                                    |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                     |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                 |

Table 142: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (Continued)

| Field Name             | Field Description                                                                                                                                  |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Type</b>            | Encapsulation being used on the interface.                                                                                                         |
| <b>Link-level type</b> | Encapsulation being used on the physical interface.                                                                                                |
| <b>MTU</b>             | MTU size on the physical interface.                                                                                                                |
| <b>Clocking</b>        | Reference clock source: can be <b>Internal</b> or <b>External</b> .                                                                                |
| <b>Speed</b>           | Speed at which the interface is running.                                                                                                           |
| <b>Device flags</b>    | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . |
| <b>Interface flags</b> | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .    |
| <b>Link type</b>       | Physical interface link type: <b>Full-Duplex</b> or <b>Half-Duplex</b> .                                                                           |
| <b>Link flags</b>      | Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .              |
| <b>Physical info</b>   | Information about the physical interface.                                                                                                          |
| <b>Hold-times</b>      | Current interface hold-time up and hold-time down, in milliseconds.                                                                                |
| <b>Current address</b> | Configured MAC address.                                                                                                                            |

Table 142: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (*Continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Alternate link address</b>  | Backup address of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <i>year-month-day hour.minute.second timezone (hour.minute.second ago)</i> . For example: <b>flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <p><b>NOTE:</b> With static NAT configured as basic NAT44 or destination NAT44 on MS-MICs and MS-MPCs, the Input bytes field might show 16 more bytes than the actual bytes in the bytes field. This is caused by the accounting of 16 bytes of the Juniper Forwarding Engine cookie.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> |

Table 142: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (Continued)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input errors</b>  | <p>Input errors on the interface. The following paragraphs explain the counters which might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame terminates and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum.</li> <li>• <b>Runts</b>—Frames received smaller than the runt threshold.</li> <li>• <b>Giants</b>—Frames received larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code does not handle because they were not recognized or not of interest. Usually, this field reports errors that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>                                              |
| <b>Output errors</b> | <p>Output errors on the interface. The following paragraphs explain the counters which might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This counter does not normally increment quickly, increasing only when the cable is unplugged from the end system is powered down and then up, or another problem occurs. If the carrier transitions increments quickly (perhaps once every 10 seconds), the cable is unplugged from the end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminates and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets larger than the MTU threshold.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> |

Table 142: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (*Continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Logical Interface</b>  |                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Logical interface</b>  | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Index</b>              | Logical interface index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>SNMP ifIndex</b>       | SNMP interface index number.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Flags</b>              | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                      |
| <b>Encapsulation</b>      | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Input packets</b>      | Number of packets received on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output packets</b>     | Number of packets transmitted on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Traffic statistics</b> | Number and rate of bytes and packets received and transmitted on the logical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> |

Table 142: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(Continued)*

| Field Name                | Field Description                                                                                                                                                                                                                                                       |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Local statistics</b>   | Statistics for traffic received from and transmitted to the Routing Engine. When traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. |
| <b>Transit statistics</b> | Statistics for traffic transiting the router. When a burst of traffic is received, the output packet rate field might briefly exceed the peak cell rate. It takes generally less than 1 second for the counter to stabilize.                                            |
| <i>protocol-family</i>    | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                |
| <b>Protocol</b>           | Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , <b>mpls</b> .                                                                                                                                                                  |
| <b>MTU</b>                | MTU size on the logical interface.                                                                                                                                                                                                                                      |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                       |
| <b>Route table</b>        | Routing table in which the logical interface address is located. For example, <b>0</b> represents the routing table <b>inet.0</b> .                                                                                                                                     |
| <b>Flags</b>              | Information about the protocol family flags. Possible values are described in the <i>Common Output Fields Description</i> section.                                                                                                                                      |
| <b>Addresses, Flags</b>   | Information about the address flags. Possible values are described in the <i>Common Output Fields Description</i> section.                                                                                                                                              |
| <b>Destination</b>        | IP address of the remote side of the connection.                                                                                                                                                                                                                        |

**Table 142: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (Continued)**

| Field Name        | Field Description                                                 |
|-------------------|-------------------------------------------------------------------|
| <b>Local</b>      | IP address of the logical interface.                              |
| <b>Broadcast</b>  | Broadcast address.                                                |
| <b>Generation</b> | Unique number for use by Juniper Networks technical support only. |

## Sample Output

### show interfaces (Adaptive Services)

```

user@host> show interfaces sp-1/2/0
Physical interface: sp-1/2/0, Enabled, Physical link is Up
 Interface index: 147, SNMP ifIndex: 72
 Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
 Speed: 800mbps
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
 Link type : Full-Duplex
 Link flags : None
 Last flapped : 2006-03-06 11:37:18 PST (00:57:29 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)

Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73)
 Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
 Input packets : 3057
 Output packets: 3044
 Protocol inet, MTU: 9192
 Flags: Receive-options, Receive-TTL-Exceeded
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.0.0.34, Local: 10.0.0.1

```

**show interfaces brief (Adaptive Services)**

```

user@host> show interfaces sp-1/2/0 brief
Physical interface: sp-1/2/0, Enabled, Physical link is Up
 Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
 Clocking: Unspecified, Speed: 800mbps
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000

Logical interface sp-1/2/0.16383
 Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
 inet 10.0.0.1 --> 10.0.0.34

```

**show interfaces detail (Adaptive Services)**

```

user@host> show interfaces sp-1/2/0 detail
Physical interface: sp-1/2/0, Enabled, Physical link is Up
 Interface index: 147, SNMP ifIndex: 72, Generation: 30
 Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
 Clocking: Unspecified, Speed: 800mbps
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
 Link type : Full-Duplex
 Link flags : None
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: Unspecified, Hardware address: Unspecified
 Alternate link address: Unspecified
 Last flapped : 2006-03-06 11:37:18 PST (00:57:56 ago)
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 125147 0 bps
 Output bytes : 1483113 0 bps
 Input packets : 3061 0 pps
 Output packets: 3048 0 pps

Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
 Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
 Traffic statistics:
 Input bytes : 125147

```

```

Output bytes : 1483113
Input packets: 3061
Output packets: 3048
Local statistics:
Input bytes : 125147
Output bytes : 1483113
Input packets: 3061
Output packets: 3048
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 9192, Generation: 20, Route table: 1
 Flags: Receive-options, Receive-TTL-Exceeded
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
 Generation: 22

```

### show interfaces extensive (Adaptive Services)

```

user@host> show interfaces sp-1/2/0 extensive
Physical interface: sp-1/2/0, Enabled, Physical link is Up
Interface index: 147, SNMP ifIndex: 72, Generation: 30
Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: 9192,
Clocking: Unspecified, Speed: 800mbps
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
Link type : Full-Duplex
Link flags : None
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Last flapped : 2006-03-06 11:37:18 PST (00:58:40 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 125547 0 bps
Output bytes : 1483353 0 bps
Input packets: 3065 0 pps
Output packets: 3052 0 pps

```

```

Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
 Policed discards: 0, Resource errors: 0
Output errors:
 Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
 Resource errors: 0

Logical interface sp-1/2/0.16383 (Index 68) (SNMP ifIndex 73) (Generation 7)
 Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
 Traffic statistics:
 Input bytes : 125547
 Output bytes : 1483353
 Input packets: 3065
 Output packets: 3052
 Local statistics:
 Input bytes : 125547
 Output bytes : 1483353
 Input packets: 3065
 Output packets: 3052
 Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
 Protocol inet, MTU: 9192, Generation: 20, Route table: 1
 Flags: Receive-options, Receive-TTL-Exceeded
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.0.0.34, Local: 10.0.0.1, Broadcast: Unspecified,
 Generation: 22

```

## Release Information

Command introduced before Junos OS Release 7.4.

## show interfaces (Aggregated Ethernet)

### IN THIS SECTION

- [Syntax | 1328](#)
- [Description | 1328](#)
- [Options | 1328](#)
- [Required Privilege Level | 1329](#)
- [Output Fields | 1329](#)
- [Sample Output | 1340](#)
- [Release Information | 1352](#)

### Syntax

```
show interfaces aenumber
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Description

(M Series, T Series, MX Series, and PTX Series routers) Display status information about the specified aggregated Ethernet interfaces.

### Options

|                                               |                                                                                                          |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b><i>aenumber</i></b>                        | Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface. |
| <b>brief   detail   extensive  <br/>terse</b> | (Optional) Display the specified level of output.                                                        |
| <b>descriptions</b>                           | (Optional) Display interface description strings.                                                        |

|                                     |                                                                               |
|-------------------------------------|-------------------------------------------------------------------------------|
| <b>media</b>                        | (Optional) Display media-specific information.                                |
| <b>snmp-index <i>snmp-index</i></b> | (Optional) Display information for the specified SNMP index of the interface. |
| <b>statistics</b>                   | (Optional) Display static interface statistics.                               |

**NOTE:** On Junos OS Evolved, in untagged aggregated ethernet (ae) interfaces with no logical interface configuration, the ae interface will not be shown as "down" and the speed will not be shown as "unspecified." The speed will be the aggregate speed of all the child member interfaces which are "up." In Junos OS, the speed is shown as "unspecified" in this case.

## Required Privilege Level

view

## Output Fields

Table 143 on page 1329 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

**Table 143: Aggregated Ethernet show interfaces Output Fields**

| Field Name                | Field Description                                                                                                                             | Level of Output |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Physical Interface</b> |                                                                                                                                               |                 |
| <b>Physical interface</b> | Name of the physical interface and state of the interface.                                                                                    | All levels      |
| <b>Enabled</b>            | State of the physical interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> . | All levels      |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                           | All levels      |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name                  | Field Description                                                                                                                                  | Level of Output                 |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>SNMP ifIndex</b>         | SNMP index number for the physical interface.                                                                                                      | <b>detail extensive</b><br>none |
| <b>Generation</b>           | Unique number for use by Juniper Networks technical support only.                                                                                  | <b>detail extensive</b>         |
| <b>Link-level type</b>      | Encapsulation being used on the physical interface.                                                                                                | All levels                      |
| <b>MTU</b>                  | Maximum transmission unit size on the physical interface.                                                                                          | All levels                      |
| <b>Speed</b>                | Speed at which the interface is running.                                                                                                           | All levels                      |
| <b>Loopback</b>             | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                     | All levels                      |
| <b>Source filtering</b>     | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                       | All levels                      |
| <b>Flow control</b>         | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                           | All levels                      |
| <b>Minimum links needed</b> | Number of child links that must be operational for the aggregate interface to be operational.                                                      | All levels                      |
| <b>Device flags</b>         | Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> . | All levels                      |
| <b>Interface flags</b>      | Information about the interface. Possible values are described in the "Interfaces Flags" section under <i>Common Output Fields Description</i> .   | All levels                      |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output         |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>Hardware address</b>        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up or from up to down. The format is <b>Last flapped: <i>year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | None specified          |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | None specified          |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes and rate, in bps, at which bytes are received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes and rate, in bps, at which bytes are transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets and rate, in pps, at which packets are received on the interface.</li> <li>• <b>Output packets</b>—Number of packets and rate, in pps, at which packets are transmitted on the interface.</li> </ul> | <b>detail extensive</b> |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output         |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Input errors</b> | <p>Input errors on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of incoming frame terminates and frame check sequence (FCS) errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's random early detection (RED) mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b> |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Output errors</b>           | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b> –Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminates and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output                 |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Queue counters</b>    | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the <b>show interfaces</b> command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p> | <b>detail extensive</b>         |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                 |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels                      |
| <b>Index</b>             | Index number of the logical interface (which reflects its initialization sequence).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b><br>none |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name           | Field Description                                                                                                                                                                                                                                                     | Level of Output                 |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>SNMP ifIndex</b>  | SNMP interface index number of the logical interface.                                                                                                                                                                                                                 | <b>detail extensive</b><br>none |
| <b>Generation</b>    | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                     | <b>detail extensive</b>         |
| <b>Flags</b>         | Information about the logical interface. Possible values are described in the “Logical Interface Flags Field” section under <i>Common Output Fields Description</i> .                                                                                                 | All levels                      |
| <b>VLAN-Tag</b>      | Tag Protocol Identifier (TPID) and VLAN identifier.                                                                                                                                                                                                                   | All levels                      |
| <b>Demux</b>         | IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: <ul style="list-style-type: none"> <li>• <b>Source Family Inet</b></li> <li>• <b>Destination Family Inet</b></li> </ul> | <b>detail extensive</b><br>none |
| <b>Encapsulation</b> | Encapsulation on the logical interface.                                                                                                                                                                                                                               | All levels                      |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Statistics</b> | <p>Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.</p> <ul style="list-style-type: none"> <li>• <b>Bundle</b>—Information about input and output bundle rates. For Junos OS Evolved, LACP packets on the members of an AE interface are not counted as part of the AE bundle input statistics.</li> <li>• <b>Link</b>—(<b>detail</b> and <b>extensive</b> only) Information about specific links in the aggregate, including link state and input and output rates.</li> <li>• <b>Adaptive Statistics</b>—(<b>extensive</b> only) Information about adaptive load balancing counter statistics. <ul style="list-style-type: none"> <li>• <b>Adaptive Adjusts</b>—Number of times traffic flow imbalance was corrected by implementation of adaptive load balancing.</li> <li>• <b>Adaptive Scans</b>—Number of times the link utilization on each member link of the AE bundle was scanned by for adaptive load balancing</li> <li>• <b>Adaptive Tolerance</b>—Tolerance level, in percentage, for load imbalance on link utilization on each member link of the AE bundle.</li> <li>• <b>Adaptive Updates</b>—Number of times traffic flow loads have been updated on an AE bundle.</li> </ul> </li> <li>• <b>Marker Statistics</b>—(<b>detail</b> and <b>extensive</b> only) Information about 802.3ad marker protocol statistics on the specified links. <ul style="list-style-type: none"> <li>• <b>Marker Rx</b>—Number of valid marker protocol data units (PDUs) received on this aggregation port.</li> </ul> </li> </ul> | <b>detail extensive</b><br>none |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <li>• <b>Resp Tx</b>—Number of marker response PDUs transmitted on this aggregation port.</li> <li>• <b>Unknown Rx</b>—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown PDU, or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.</li> <li>• <b>Illegal Rx</b>—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).</li> </ul> |                 |

Table 143: Aggregated Ethernet show interfaces Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| LACP info  | <p>Link Aggregation Control Protocol (LACP) information for each aggregated interface.</p> <ul style="list-style-type: none"> <li>• <b>Role</b> can be one of the following: <ul style="list-style-type: none"> <li>• <b>Actor</b>—Local device participating in LACP negotiation.</li> <li>• <b>Partner</b>—Remote device participating in LACP negotiation.</li> </ul> </li> <li>• <b>System priority</b>—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>System identifier</b>—Actor or partner system ID, encoded as a MAC address.</li> <li>• <b>Port priority</b>—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>Unknown Rx</b>—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown protocol data unit (PDU), or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.</li> <li>• <b>Port key</b>—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.</li> </ul> |                 |

Table 143: Aggregated Ethernet show interfaces Output Fields (*Continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output                 |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| LACP Statistics        | <p>LACP statistics for each aggregated interface.</p> <ul style="list-style-type: none"> <li>• <b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li>• <b>LACP Tx</b>—Number of LACP transmit packet errors logged.</li> <li>• <b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li>• <b>Illegal Rx</b>—Number of invalid packets received.</li> </ul> <p><b>NOTE:</b> For <b>LACP Rx</b> and <b>LACP Tx</b>, Packet count is updated only on snmp timer expiry (30 secs).</p> | <b>detail extensive</b><br>none |
| <i>protocol-family</i> | Protocol family configured on the logical interface. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                               | <b>brief</b>                    |
| <b>Protocol</b>        | Protocol family configured on the logical interface. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b><br>none |
| <b>MTU</b>             | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b><br>none |
| <b>Maximum labels</b>  | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b><br>none |
| <b>Generation</b>      | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>         |
| <b>Route Table</b>     | Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |

Table 143: Aggregated Ethernet show interfaces Output Fields (*Continued*)

| Field Name                   | Field Description                                                                                                                                          | Level of Output                 |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Flags</b>                 | Information about protocol family flags. Possible values are described in the “Family Flags Field” section under <i>Common Output Fields Description</i> . | <b>detail extensive</b><br>none |
| <b>Mac-Validate Failures</b> | Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface. | <b>detail extensive</b><br>none |
| <b>Addresses, Flags</b>      | Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .            | <b>detail extensive</b><br>none |
| <b>Destination</b>           | IP address of the remote side of the connection.                                                                                                           | <b>detail extensive</b><br>none |
| <b>Local</b>                 | IP address of the logical interface.                                                                                                                       | <b>detail extensive</b><br>none |
| <b>Broadcast</b>             | Broadcast address of the logical interface.                                                                                                                | <b>detail extensive</b><br>none |
| <b>Generation</b>            | Unique number for use by Juniper Networks technical support only.                                                                                          | <b>detail extensive</b>         |

## Sample Output

### show interfaces (Aggregated Ethernet)

```

user@host> show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
 Interface index: 153, SNMP ifIndex: 59
 Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,

```

```

Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Current address: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
Last flapped : Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

```

```

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60)

```

```

Flags: SNMP-Traps 16384 Encapsulation: ENET2

```

| Statistics | Packets | pps | Bytes | bps |
|------------|---------|-----|-------|-----|
| Bundle:    |         |     |       |     |
| Input :    | 0       | 0   | 0     | 0   |
| Output:    | 0       | 0   | 0     | 0   |

```

Protocol inet, MTU: 1500

```

```

Flags: None

```

```

Addresses, Flags: Is-Preferred Is-Primary

```

```

Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 10.100.1.255

```

### show interfaces brief (Aggregated Ethernet)

```

user@host> show interfaces ae0 brief
Physical interface: ae0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps 16384

Logical interface ae0.0
Flags: SNMP-Traps 16384 Encapsulation: ENET2
inet 203.0.113.2/24

```

### show interfaces detail (Aggregated Ethernet)

```

user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
Interface index: 153, SNMP ifIndex: 59, Generation: 36
Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1

```

```

Device flags : Present Running
Interface flags: SNMP-Traps 16384
Current address: 00:00:5e:00:53:f0, Hardware address: 00:00:5e:00:53:f0
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Queue counters: Queued packets Transmitted packets Dropped packets
 0 best-effort 7375 7375 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont 2268 2268 0

Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
Flags: SNMP-Traps 16384 Encapsulation: ENET2
Statistics Packets pps Bytes bps
Bundle:
 Input : 0 0 0 0
 Output: 0 0 0 0
Link:
 fe-0/1/0.0
 Input : 0 0 0 0
 Output: 0 0 0 0
 fe-0/1/2.0
 Input : 0 0 0 0
 Output: 0 0 0 0
 fe-0/1/3.0
 Input : 0 0 0 0
 Output: 0 0 0 0
Marker Statistics: Marker Rx Resp Tx Unknown Rx Illegal Rx
 fe-0/1/0.0 0 0 0 0
 fe-0/1/2.0 0 0 0 0
 fe-0/1/3.0 0 0 0 0
Protocol inet, MTU: 1500, Generation: 37, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
 Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
 Generation: 49

```

**show interfaces extensive (Aggregated Ethernet)**

```

user@host> show interfaces ae0 extensive
Physical interface: ae1, Enabled, Physical link is Up
 Interface index: 304, SNMP ifIndex: 830, Generation: 342
 Link-level type: Flexible-Ethernet, MTU: 16000, Speed: 200Gbps, 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: 00:00:00:ab:cd:e1, Hardware address: f6:cc:55:34:af:8f
 Last flapped : 2019-11-04 14:09:06 IST (00:21:46 ago)
 Statistics last cleared: 2019-11-04 14:30:46 IST (00:00:06 ago)
 Traffic statistics:
 Input bytes : 12955595850 16067340264 bps
 Output bytes : 12955839144 16067696488 bps
 Input packets: 12904045 2000425 pps
 Output packets: 12904281 2000469 pps
 IPv6 transit statistics:
 Input bytes : 12955586722
 Output bytes : 12955830836
 Input packets: 12903971
 Output packets: 12904214
 Dropped traffic statistics due to STP State:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 MAC statistics:
 Broadcast packets Receive Transmit
 Multicast packets 11 11
 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
 Egress queues: 8 supported, 4 in use
 Queue counters:
 Queued packets Transmitted packets Dropped packets
 0 12623539 12623539 0

```

```

1 0 0 0
2 0 0 0
3 67 67 0
Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control

Logical interface ae1.1501 (Index 405) (SNMP ifIndex 982) (Generation 12278)
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.1501] Encapsulation: ENET2
Statistics Packets pps Bytes bps
Bundle:
 Input : 12903972 2000413 12955586798 16067363080
 Output: 12904214 2000456 12955830836 16067715624
Adaptive Statistics:
 Adaptive Adjusts: 0
 Adaptive Scans : 0
 Adaptive Updates: 0
Link:
 et-8/7/4.1501
 Input : 6453558 999544 6479365950 8028338176
 Output: 5419689 840489 5441367750 6750815312
 xe-8/0/0:0.1501
 Input : 630823 97367 633346374 782052320
 Output: 946439 146652 950224754 1177910328
 xe-8/0/0:1.1501
 Input : 681646 105969 684374260 851151032
 Output: 860269 133412 863710070 1071571528
 xe-8/0/0:2.1501
 Input : 629697 97558 632215406 783587288
 Output: 774250 120067 777347000 964383592
 xe-8/0/0:3.1501
 Input : 630012 98027 632532324 787360080
 Output: 1032345 160054 1036474378 1285560688
 xe-8/6/0:0.1501
 Input : 654743 101651 657362010 816468784
 Output: 473122 73326 475100508 589061512
 xe-8/6/0:1.1501
 Input : 630360 97905 632883672 786375112
 Output: 860240 133284 863680960 1070540592
 xe-8/6/0:2.1501
 Input : 654850 101450 657469796 814853936

```

```

Output: 688226 106642 690978904 856554888
xe-8/7/0:1.1501
Input : 654370 101442 656986660 814784400
Output: 516173 79971 518237690 642332976
xe-8/7/0:2.1501
Input : 655771 101940 658394074 818781864
Output: 1075371 166587 1079586458 1337921528
xe-8/7/0:3.1501
Input : 628142 97560 630656272 783610088
Output: 258090 39972 259122364 321062680

```

Aggregate member links: 11

| Marker Statistics: | Marker Rx | Resp Tx | Unknown Rx | Illegal Rx |
|--------------------|-----------|---------|------------|------------|
| et-8/7/4.1501      | 0         | 0       | 0          | 0          |
| xe-8/0/0:0.1501    | 0         | 0       | 0          | 0          |
| xe-8/0/0:1.1501    | 0         | 0       | 0          | 0          |
| xe-8/0/0:2.1501    | 0         | 0       | 0          | 0          |
| xe-8/0/0:3.1501    | 0         | 0       | 0          | 0          |
| xe-8/6/0:0.1501    | 0         | 0       | 0          | 0          |
| xe-8/6/0:1.1501    | 0         | 0       | 0          | 0          |
| xe-8/6/0:2.1501    | 0         | 0       | 0          | 0          |
| xe-8/7/0:1.1501    | 0         | 0       | 0          | 0          |
| xe-8/7/0:2.1501    | 0         | 0       | 0          | 0          |
| xe-8/7/0:3.1501    | 0         | 0       | 0          | 0          |

Protocol inet, MTU: 15978

Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop cnt: 0

Generation: 36373, Route table: 0

Flags: Sendbcast-pkt-to-re

Addresses, Flags: Is-Preferred Is-Primary

Destination: 172.161.1/24, Local: 172.161.1.1, Broadcast: 172.161.1.255,

Generation: 72255

Protocol inet6, MTU: 15978

Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 2, Curr new hold cnt: 0, NH drop cnt: 0

Generation: 36374, Route table: 0

Addresses, Flags: Is-Preferred Is-Primary

Destination: 2005:1:161:1::/64, Local: 2005:1:161:1::1

Generation: 72257

Addresses, Flags: Is-Preferred

Destination: fe80::/64, Local: fe80::f6cc:5505:dd34:af8f

Protocol multiservice, MTU: Unlimited, Generation: 72259

Generation: 36375, Route table: 0

Policer: Input: \_\_default\_arp\_policer\_\_

Logical interface ae1.32767 (Index 406) (SNMP ifIndex 985) (Generation 12279)

Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

| Statistics | Packets | pps | Bytes | bps |
|------------|---------|-----|-------|-----|
|------------|---------|-----|-------|-----|

Bundle:

|         |    |   |      |       |
|---------|----|---|------|-------|
| Input : | 73 | 1 | 9052 | 10488 |
|---------|----|---|------|-------|

|         |    |   |      |       |
|---------|----|---|------|-------|
| Output: | 67 | 0 | 8308 | 10032 |
|---------|----|---|------|-------|

Adaptive Statistics:

|                   |   |
|-------------------|---|
| Adaptive Adjusts: | 0 |
|-------------------|---|

|                  |   |
|------------------|---|
| Adaptive Scans : | 0 |
|------------------|---|

|                   |   |
|-------------------|---|
| Adaptive Updates: | 0 |
|-------------------|---|

Link:

et-8/7/4.32767

|         |   |   |     |      |
|---------|---|---|-----|------|
| Input : | 7 | 1 | 868 | 1368 |
|---------|---|---|-----|------|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 7 | 0 | 868 | 912 |
|---------|---|---|-----|-----|

xe-8/0/0:0.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/0/0:1.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/0/0:2.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 7 | 0 | 868 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/0/0:3.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 7 | 0 | 868 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/6/0:0.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 7 | 0 | 868 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/6/0:1.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 7 | 0 | 868 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/6/0:2.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 7 | 0 | 868 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/7/0:1.32767

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Input : | 7 | 0 | 868 | 912 |
|---------|---|---|-----|-----|

|         |   |   |     |     |
|---------|---|---|-----|-----|
| Output: | 6 | 0 | 744 | 912 |
|---------|---|---|-----|-----|

xe-8/7/0:2.32767

|      |                  |         |          |                   |          |        |
|------|------------------|---------|----------|-------------------|----------|--------|
|      | Input :          | 6       | 0        | 744               | 912      |        |
|      | Output:          | 6       | 0        | 744               | 912      |        |
|      | xe-8/7/0:3.32767 |         |          |                   |          |        |
|      | Input :          | 6       | 0        | 744               | 912      |        |
|      | Output:          | 6       | 0        | 744               | 912      |        |
|      | LACP info:       | Role    | System   | System            | Port     | Port   |
| Port |                  |         | priority | identifier        | priority | number |
| key  |                  |         |          |                   |          |        |
| 26   | et-8/7/4.32767   | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 8    | et-8/7/4.32767   | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 3    | xe-8/0/0:0.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 16   | xe-8/0/0:0.32767 | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 4    | xe-8/0/0:1.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 17   | xe-8/0/0:1.32767 | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 5    | xe-8/0/0:2.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 18   | xe-8/0/0:2.32767 | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 6    | xe-8/0/0:3.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 19   | xe-8/0/0:3.32767 | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 17   | xe-8/6/0:0.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 12   | xe-8/6/0:0.32767 | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 18   | xe-8/6/0:1.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 13   | xe-8/6/0:1.32767 | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 19   | xe-8/6/0:2.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 14   | xe-8/6/0:2.32767 | Partner | 127      | d8:b1:22:f2:57:c0 | 127      |        |
|      | 2                |         |          |                   |          |        |
| 22   | xe-8/7/0:1.32767 | Actor   | 127      | f4:cc:55:34:b7:40 | 127      |        |
|      | 2                |         |          |                   |          |        |

|    |                  |         |     |                   |     |
|----|------------------|---------|-----|-------------------|-----|
| 20 | xe-8/7/0:1.32767 | Partner | 127 | d8:b1:22:f2:57:c0 | 127 |
|    | 2                |         |     |                   |     |
| 23 | xe-8/7/0:2.32767 | Actor   | 127 | f4:cc:55:34:b7:40 | 127 |
|    | 2                |         |     |                   |     |
| 21 | xe-8/7/0:2.32767 | Partner | 127 | d8:b1:22:f2:57:c0 | 127 |
|    | 2                |         |     |                   |     |
| 24 | xe-8/7/0:3.32767 | Actor   | 127 | f4:cc:55:34:b7:40 | 127 |
|    | 2                |         |     |                   |     |
| 22 | xe-8/7/0:3.32767 | Partner | 127 | d8:b1:22:f2:57:c0 | 127 |
|    | 2                |         |     |                   |     |

| LACP Statistics: | LACP Rx | LACP Tx | Unknown Rx | Illegal Rx |
|------------------|---------|---------|------------|------------|
| et-8/7/4.32767   | 0       | 0       | 0          | 0          |
| xe-8/0/0:0.32767 | 0       | 0       | 0          | 0          |
| xe-8/0/0:1.32767 | 0       | 0       | 0          | 0          |
| xe-8/0/0:2.32767 | 0       | 0       | 0          | 0          |
| xe-8/0/0:3.32767 | 0       | 0       | 0          | 0          |
| xe-8/6/0:0.32767 | 0       | 0       | 0          | 0          |
| xe-8/6/0:1.32767 | 0       | 0       | 0          | 0          |
| xe-8/6/0:2.32767 | 0       | 0       | 0          | 0          |
| xe-8/7/0:1.32767 | 0       | 0       | 0          | 0          |
| xe-8/7/0:2.32767 | 0       | 0       | 0          | 0          |
| xe-8/7/0:3.32767 | 0       | 0       | 0          | 0          |

| Marker Statistics: | Marker Rx | Resp Tx | Unknown Rx | Illegal Rx |
|--------------------|-----------|---------|------------|------------|
| et-8/7/4.32767     | 0         | 0       | 0          | 0          |
| xe-8/0/0:0.32767   | 0         | 0       | 0          | 0          |
| xe-8/0/0:1.32767   | 0         | 0       | 0          | 0          |
| xe-8/0/0:2.32767   | 0         | 0       | 0          | 0          |
| xe-8/0/0:3.32767   | 0         | 0       | 0          | 0          |
| xe-8/6/0:0.32767   | 0         | 0       | 0          | 0          |
| xe-8/6/0:1.32767   | 0         | 0       | 0          | 0          |
| xe-8/6/0:2.32767   | 0         | 0       | 0          | 0          |
| xe-8/7/0:1.32767   | 0         | 0       | 0          | 0          |
| xe-8/7/0:2.32767   | 0         | 0       | 0          | 0          |
| xe-8/7/0:3.32767   | 0         | 0       | 0          | 0          |

Protocol multiservice, MTU: Unlimited, Generation: 36376, Route table: 0

Flags: None

Policer: Input: \_\_default\_arp\_policer\_\_

**show interfaces extensive (Aggregated Ethernet with VLAN Stacking)**

```

user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
 Interface index: 155, SNMP ifIndex: 48, Generation: 186
 Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source
 filtering: Disabled,
 Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Current address: 00:00:5e:00:53:3f, Hardware address: 00:00:5e:00:53:3f
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 2406875 40152 bps
 Output bytes : 1124470 22056 bps
 Input packets : 5307 5 pps
 Output packets: 13295 21 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, 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: 4 supported, 4 in use
 Queue counters:
 Queued packets Transmitted packets Dropped packets
 0 best-effort 0 859777 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont 0 0 0
 Egress queues: 4 supported, 4 in use
 Queue counters:
 Queued packets Transmitted packets Dropped packets
 0 best-effort 0 1897615 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont 0 662505 0

```

Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)

Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ] Encapsulation: VLAN-VPLS

| Statistics | Packets | pps | Bytes | bps |
|------------|---------|-----|-------|-----|
|------------|---------|-----|-------|-----|

Bundle:

|         |     |   |       |     |
|---------|-----|---|-------|-----|
| Input : | 289 | 0 | 25685 | 376 |
|---------|-----|---|-------|-----|

|         |      |   |        |      |
|---------|------|---|--------|------|
| Output: | 1698 | 4 | 130375 | 3096 |
|---------|------|---|--------|------|

Link:

ge-1/2/0.451

|         |     |   |       |     |
|---------|-----|---|-------|-----|
| Input : | 289 | 0 | 25685 | 376 |
|---------|-----|---|-------|-----|

|         |   |   |   |   |
|---------|---|---|---|---|
| Output: | 0 | 0 | 0 | 0 |
|---------|---|---|---|---|

ge-1/2/1.451

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
|---------|---|---|---|---|

|         |      |   |        |      |
|---------|------|---|--------|------|
| Output: | 1698 | 4 | 130375 | 3096 |
|---------|------|---|--------|------|

| Marker Statistics: | Marker Rx | Resp Tx | Unknown Rx | Illegal Rx |
|--------------------|-----------|---------|------------|------------|
|--------------------|-----------|---------|------------|------------|

|              |   |   |   |   |
|--------------|---|---|---|---|
| ge-1/2/0.451 | 0 | 0 | 0 | 0 |
|--------------|---|---|---|---|

|              |   |   |   |   |
|--------------|---|---|---|---|
| ge-1/2/1.451 | 0 | 0 | 0 | 0 |
|--------------|---|---|---|---|

Protocol vpls, MTU: 1518, Generation: 849, Route table: 3

Flags: Is-Primary

Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)

Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ] Encapsulation: VLAN-VPLS

| Statistics | Packets | pps | Bytes | bps |
|------------|---------|-----|-------|-----|
|------------|---------|-----|-------|-----|

Bundle:

|         |     |   |       |      |
|---------|-----|---|-------|------|
| Input : | 293 | 1 | 26003 | 1072 |
|---------|-----|---|-------|------|

|         |      |   |        |      |
|---------|------|---|--------|------|
| Output: | 1694 | 3 | 130057 | 2400 |
|---------|------|---|--------|------|

Link:

ge-1/2/0.452

|         |     |   |       |      |
|---------|-----|---|-------|------|
| Input : | 293 | 1 | 26003 | 1072 |
|---------|-----|---|-------|------|

|         |      |   |        |      |
|---------|------|---|--------|------|
| Output: | 1694 | 3 | 130057 | 2400 |
|---------|------|---|--------|------|

ge-1/2/1.452

|         |   |   |   |   |
|---------|---|---|---|---|
| Input : | 0 | 0 | 0 | 0 |
|---------|---|---|---|---|

|         |   |   |   |   |
|---------|---|---|---|---|
| Output: | 0 | 0 | 0 | 0 |
|---------|---|---|---|---|

| Marker Statistics: | Marker Rx | Resp Tx | Unknown Rx | Illegal Rx |
|--------------------|-----------|---------|------------|------------|
|--------------------|-----------|---------|------------|------------|

|              |   |   |   |   |
|--------------|---|---|---|---|
| ge-1/2/0.452 | 0 | 0 | 0 | 0 |
|--------------|---|---|---|---|

|              |   |   |   |   |
|--------------|---|---|---|---|
| ge-1/2/1.452 | 0 | 0 | 0 | 0 |
|--------------|---|---|---|---|

Protocol vpls, MTU: 1518, Generation: 850, Route table: 3

Flags: None

...

**show interfaces extensive (Aggregated Ethernet for PTX100008)**

```

user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Down
 Interface index: 1013, SNMP ifIndex: 513, Generation: 670014899862
 Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 1bps,

Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps
Current address: 20:67:c4:17:05:d5, Hardware address: 20:67:c4:17:05:d5
Last flapped : 2020-09-01 17:41:00 IST (00:03:35 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Police
discrads: 0, Resource errors: 0
 Output errors:

 Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
 Adaptive Statistics:
 Adaptive Adjusts: 30
 Adaptive Scans : 10
 Adaptive Updates: 30

Logical interface ae0.0 (Index 1003) (SNMP ifIndex 514) (Generation
670014899865)
 Flags: Up SNMP-Traps 16384 Encapsulation: ENET2
 Statistics Packets pps Bytes bps
 Bundle:
 Input : 0 0 0 0
 Output : 0 0 0 0
 Aggregate member links: 0
 Protocol inet, MTU: 1500
 Generation: : 670014899866, Route table: 0, Next-hop: 0
 Flags: Sendbcast-pkt-to-re, Is-Primary

```

```
Addresses, Flags: Is-Preferred Is-Primary
Destination: Unspecified, Local: 10.10.10.1, Broadcast: Unspecified,
Generation: 670014899868
Protocol multiservice, MTU: Unlimited, Generation: 670014899871, Route
table: 0, Next-hop: 0
```

Starting in Junos OS Evolved Release 20.3R1, when you configure adaptive load balancing, the `show interfaces ae0 extensive` command displays adaptive statistics under the physical interface, not the logical interface for the PTX10008 Series routers.

## Release Information

Command introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [Ethernet Interfaces User Guide for Routing Devices](#)

## show interfaces demux0 (Demux Interfaces)

### IN THIS SECTION

- [Syntax | 1353](#)
- [Description | 1353](#)
- [Options | 1353](#)
- [Required Privilege Level | 1353](#)
- [Output Fields | 1353](#)
- [Sample Output | 1365](#)
- [Release Information | 1368](#)

## Syntax

```
show interfaces demux0.logical-interface-number
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Description

(MX Series and M Series routers only) Display status information about the specified demux interface.

## Options

|                                           |                                                                               |
|-------------------------------------------|-------------------------------------------------------------------------------|
| <b>none</b>                               | Display standard information about the specified demux interface.             |
| <b>brief   detail   extensive   terse</b> | (Optional) Display the specified level of output.                             |
| <b>descriptions</b>                       | (Optional) Display interface description strings.                             |
| <b>media</b>                              | (Optional) Display media-specific information about network interfaces.       |
| <b>snmp-index <i>snmp-index</i></b>       | (Optional) Display information for the specified SNMP index of the interface. |
| <b>statistics</b>                         | (Optional) Display static interface statistics.                               |

## Required Privilege Level

view

## Output Fields

[Table 144 on page 1354](#) lists the output fields for the **show interfaces demux0** (Demux Interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields

| Field Name                | Field Description                                                                                                                             | Level of Output                              |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <b>Physical Interface</b> |                                                                                                                                               |                                              |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                               | <b>brief detail</b><br><b>extensive none</b> |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                           | <b>brief detail</b><br><b>extensive none</b> |
| <b>Enabled</b>            | State of the interface. Possible values are described in the "Enabled Field" section under <a href="#">Common Output Fields Description</a> . | <b>brief detail</b><br><b>extensive none</b> |
| <b>Physical link</b>      | Status of the physical link ( <b>Up</b> or <b>Down</b> ).                                                                                     | <b>detail extensive</b><br><b>none</b>       |
| <b>Admin</b>              | Administrative state of the interface ( <b>Up</b> or <b>Down</b> ).                                                                           | <b>terse</b>                                 |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                           | <b>detail extensive</b><br><b>none</b>       |
| <b>Link</b>               | Status of the physical link ( <b>Up</b> or <b>Down</b> ).                                                                                     | <b>terse</b>                                 |
| <b>Targeting summary</b>  | Status of aggregated Ethernet links that are configured with targeted distribution ( <b>primary</b> or <b>backup</b> )                        | <b>extensive</b>                             |
| <b>Bandwidth</b>          | Bandwidth allocated to the aggregated Ethernet links that are configured with targeted distribution.                                          | <b>extensive</b>                             |
| <b>Proto</b>              | Protocol family configured on the interface.                                                                                                  | <b>terse</b>                                 |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name             | Field Description                                                                                                                                           | Level of Output                              |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <b>SNMP ifIndex</b>    | SNMP index number for the physical interface.                                                                                                               | <b>detail extensive</b><br>none              |
| <b>Generation</b>      | Unique number for use by Juniper Networks technical support only.                                                                                           | <b>detail extensive</b>                      |
| <b>Type</b>            | Type of interface. <b>Software-Pseudo</b> indicates a standard software interface with no associated hardware device.                                       | <b>brief detail</b><br><b>extensive</b> none |
| <b>Link-level type</b> | Encapsulation being used on the physical interface.                                                                                                         | <b>brief detail</b><br><b>extensive</b>      |
| <b>MTU</b>             | Maximum transmission unit size on the physical interface.                                                                                                   | <b>brief detail</b><br><b>extensive</b>      |
| <b>Clocking</b>        | Reference clock source: <b>Internal</b> (1) or <b>External</b> (2).                                                                                         | <b>brief detail</b><br><b>extensive</b>      |
| <b>Speed</b>           | Speed at which the interface is running.                                                                                                                    | <b>brief detail</b><br><b>extensive</b>      |
| <b>Device flags</b>    | Information about the physical device. Possible values are described in the "Device Flags" section under <a href="#">Common Output Fields Description</a> . | <b>brief detail</b><br><b>extensive</b> none |
| <b>Interface flags</b> | Information about the interface. Possible values are described in the "Interface Flags" section under <a href="#">Common Output Fields Description</a> .    | <b>brief detail</b><br><b>extensive</b> none |
| <b>Link type</b>       | Data transmission type.                                                                                                                                     | <b>detail extensive</b><br>none              |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                           | Level of Output                 |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Link flags</b>              | Information about the link. Possible values are described in the “Link Flags” section under <a href="#">Common Output Fields Description</a> .                                                                                                              | <b>detail extensive</b><br>none |
| <b>Physical info</b>           | Information about the physical interface.                                                                                                                                                                                                                   | <b>detail extensive</b>         |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                         | <b>detail extensive</b>         |
| <b>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                     | <b>detail extensive</b>         |
| <b>Hardware address</b>        | Hardware MAC address.                                                                                                                                                                                                                                       | <b>detail extensive</b>         |
| <b>Alternate link address</b>  | Backup address of the link.                                                                                                                                                                                                                                 | <b>detail extensive</b>         |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive</b><br>none |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                           | <b>detail extensive</b>         |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output         |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> <li>• <b>IPv6 transit statistics</b>—Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</li> </ul> <p><b>NOTE:</b> These fields include dropped traffic and exception traffic, as those fields are not separately defined.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b> |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output  |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b> | <p>Input errors on the interface whose definitions are as follows:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame terminations and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant packet threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |
| <b>Input Rate</b>   | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | none             |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output                              |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <b>Output errors</b>     | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminations and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>                             |
| <b>Output Rate</b>       | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | none                                         |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                              |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>brief detail</b><br><b>extensive</b> none |
| <b>Index</b>             | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b><br>none              |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output                              |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <b>SNMP ifIndex</b>              | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b><br>none              |
| <b>Generation</b>                | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>                                |
| <b>Flags</b>                     | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <a href="#">Common Output Fields Description</a> .                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>brief detail</b><br><b>extensive</b> none |
| <b>Encapsulation</b>             | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>brief extensive</b><br>none               |
| <b>ACI VLAN: Dynamic Profile</b> | Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying demux interface to create dynamic VLAN subscriber interfaces based on ACI information.                                                                                                                                                                                                                                                                                                                                                        | <b>brief detail</b><br><b>extensive</b> none |
| <b>Demux</b>                     | Specific IP demultiplexing (demux) values: <ul style="list-style-type: none"> <li>• <b>Underlying interface</b>—The underlying interface that the demux interface uses.</li> <li>• <b>Index</b>—Index number of the logical interface.</li> <li>• <b>Family</b>—Protocol family configured on the logical interface.</li> <li>• <b>Source prefixes, total</b>—Total number of source prefixes for the underlying interface.</li> <li>• <b>Destination prefixes, total</b>—Total number of destination prefixes for the underlying interface.</li> <li>• <b>Prefix</b>—inet family prefix.</li> </ul> | <b>detail extensive</b><br>none              |
| <i>protocol-family</i>           | Protocol family configured on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>brief</b>                                 |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output         |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> <li>• <b>IPv6 transit statistics</b>—Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</li> </ul> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b> |
| <b>Local statistics</b>   | <p>Number of transit bytes and packets received and transmitted on the local interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b> |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output                 |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Transit statistics</b>      | <p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                        | <b>detail extensive</b>         |
| <b>IPv6 Transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |
| <b>Input packets</b>           | Number of packets received on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | none                            |
| <b>Output packets</b>          | Number of packets transmitted on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | none                            |
| <b>Protocol</b>                | Protocol family. Possible values are described in the “Protocol Field” section under <a href="#">Common Output Fields Description</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b><br>none |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name                   | Field Description                                                                                                                                             | Level of Output                                   |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| <b>MTU</b>                   | Maximum transmission unit size on the logical interface.                                                                                                      | <b>detail extensive</b><br>none                   |
| <b>Maximum labels</b>        | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                               | <b>detail extensive</b><br>none                   |
| <b>Generation</b>            | Unique number for use by Juniper Networks technical support only.                                                                                             | <b>detail extensive</b>                           |
| <b>Route table</b>           | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                      | <b>detail extensive</b>                           |
| <b>Flags</b>                 | Information about protocol family flags. Possible values are described in the “Family Flags” section under <a href="#">Common Output Fields Description</a> . | <b>detail extensive</b><br>none                   |
| <b>Mac-Validate Failures</b> | Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.    | <b>detail extensive</b><br>none                   |
| <b>Addresses, Flags</b>      | Information about the address flags. Possible values are described in the “Addresses Flags” section under <a href="#">Common Output Fields Description</a> .  | <b>detail extensive</b><br>none                   |
| <b>Destination</b>           | IP address of the remote side of the connection.                                                                                                              | <b>detail extensive</b><br><b>statistics</b> none |
| <b>Local</b>                 | IP address of the logical interface.                                                                                                                          | <b>detail extensive</b><br><b>terse</b> none      |
| <b>Remote</b>                | IP address of the remote interface.                                                                                                                           | <b>terse</b>                                      |

Table 144: show interfaces demux0 (Demux Interfaces) Output Fields (Continued)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                    | Level of Output                 |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Broadcast</b>            | Broadcast address of the logical interlace.                                                                                                                                                                                                                                                          | <b>detail extensive</b><br>none |
| <b>Generation</b>           | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                    | <b>detail extensive</b>         |
| <b>Link</b>                 | Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.                                                                                                                    | <b>detail extensive</b><br>none |
| <b>Dynamic-profile</b>      | Name of the PPPoE dynamic profile assigned to the underlying interface.                                                                                                                                                                                                                              | <b>detail extensive</b><br>none |
| <b>Service Name Table</b>   | Name of the PPPoE service name table assigned to the PPPoE underlying interface.                                                                                                                                                                                                                     | <b>detail extensive</b><br>none |
| <b>Max Sessions</b>         | Maximum number of dynamic PPPoE logical interfaces that the router can activate on the underlying interface.                                                                                                                                                                                         | <b>detail extensive</b><br>none |
| <b>Duplicate Protection</b> | State of duplicate protection: <b>On</b> or <b>Off</b> . Duplicate protection prevents the activation of another dynamic PPPoE logical interface on the same underlying interface when a dynamic PPPoE logical interface for a client with the same MAC address is already active on that interface. | <b>detail extensive</b><br>none |
| <b>Direct Connect</b>       | State of the configuration to ignore DSL Forum VSAs: <b>On</b> or <b>Off</b> . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.                                                                                                 | <b>detail extensive</b><br>none |
| <b>AC Name</b>              | Name of the access concentrator.                                                                                                                                                                                                                                                                     | <b>detail extensive</b><br>none |

## Sample Output

### show interfaces demux0 (Demux)

```

user@host> show interfaces demux0
Physical interface: demux0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 79, Generation: 129
 Type: Software-Pseudo, Link-level type: Unspecified, MTU: 9192, Clocking: 1,
 Speed: Unspecified
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps
 Link type : Full-Duplex
 Link flags : None
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: Unspecified, Hardware address: Unspecified
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps
 Output packets: 0 0 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
 Policed discards: 0, Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
 Resource errors: 0

Logical interface demux0.0 (Index 87) (SNMP ifIndex 84) (Generation 312)
 Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
 Demux:
 Underlying interface: ge-2/0/1.0 (Index 74)
 Family Inet Source prefixes, total 1
 Prefix: 203.0.113/24

```

```

Traffic statistics:
 Input bytes : 0
 Output bytes : 1554
 Input packets: 0
 Output packets: 37
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 1554
 Input packets: 0
 Output packets: 37
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 395, Route table: 0
 Flags: Is-Primary, Mac-Validate-Strict
 Mac-Validate Failures: Packets: 0, Bytes: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 203.0.113/24, Local: 203.0.113.13, Broadcast: 203.0.113.255,
 Generation: 434

```

### show interfaces demux0 (PPPoE over Aggregated Ethernet)

```

user@host> show interfaces demux0.100
Logical interface demux0.100 (Index 76) (SNMP ifIndex 61160)
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.100]
 Encapsulation: ENET2
 Demux:
 Underlying interface: ae0 (Index 199)
 Link:

```

```

 ge-1/0/0
 ge-1/1/0
Input packets : 0
Output packets: 0
Protocol pppoe
 Dynamic Profile: pppoe-profile,
 Service Name Table: service-table1,
 Max Sessions: 100, Duplicate Protection: On,
 Direct Connect: Off,
 AC Name: pppoe-server-1

```

### show interfaces demux0 extensive (Targeted Distribution for Aggregated Ethernet Links)

```

user@host> show interfaces demux0.1073741824 extensive

Logical interface demux0.1073741824 (Index 75) (SNMP ifIndex 558) (Generation
346)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
Demux:
 Underlying interface: ae0 (Index 201)
Link:
 ge-1/0/0
 ge-1/1/0
 ge-2/0/7
 ge-2/0/8
Targeting summary:
 ge-1/1/0, primary, Physical link is Up
 ge-2/0/8, backup, Physical link is Up
Bandwidth: 1000mbps

```

### show interfaces demux0 (ACI Interface Set Configured)

```

user@host> show interfaces demux0.1073741827

Logical interface demux0.1073741827 (Index 346) (SNMP ifIndex 527)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1802 0x8100.302] Encapsulation:
ENET2
Demux: Source Family Inet
ACI VLAN:
 Dynamic Profile: aci-vlan-set-profile

```

```

Demux:
 Underlying interface: ge-1/0/0 (Index 138)
Input packets : 18
Output packets: 16
Protocol inet, MTU: 1500
 Flags: Sendbcast-pkt-to-re, Unnumbered
 Donor interface: lo0.0 (Index 322)
 Preferred source address: 203.0.113.202
 Addresses, Flags: Primary Is-Default Is-Primary
 Local: 203.0.113.119
Protocol pppoe
 Dynamic Profile: aci-vlan-pppoe-profile,
 Service Name Table: None,
 Max Sessions: 32000, Max Sessions VSA Ignore: Off,
 Duplicate Protection: On, Short Cycle Protection: Off,
 Direct Connect: Off,
 AC Name: nbc

```

## Release Information

Command introduced in Junos OS Release 9.0.

## RELATED DOCUMENTATION

| [Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration](#)

## show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port)

### IN THIS SECTION

- [Syntax | 1369](#)
- [Description | 1369](#)
- [Options | 1369](#)
- [Additional Information | 1369](#)

- Required Privilege Level | 1370
- Output Fields | 1370
- Sample Output | 1403
- Release Information | 1431

## Syntax

```
show interfaces diagnostics optics interface-name
```

## Description

Display diagnostics data, warnings, and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, or Virtual Chassis port interfaces.

## Options

*interface-name*

Interface name. For example:

*ge-fpc/pic/port*

*et-fpc/pic/port*

*et-fpc/pic/port.channel*

*xe-fpc/pic/port*

*vcp-fpc/pic/port*

## Additional Information

The transceivers are polled in 1-second intervals for diagnostics data, warnings, and alarms. The alarms do not cause the links to go down or the LEDs to change color, nor generate SNMP traps. Changes in alarm and warning status generate system log messages.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transceiver vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a device is not working.

**NOTE:** Some transceivers do not support all optical diagnostics features described in the output fields.

If optics measures transmit or receive power as zero, then, the measured power is displayed as **0.000 mW / - Inf dBm**

You can configure the P2-10G-40G-QSFPP PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode. When the PIC is in 40-Gigabit Ethernet mode, you must execute the **show interfaces diagnostics optics et-fpcl/picl/port** command. The output of this command displays the diagnostic optics information about the corresponding 40-Gigabit Ethernet port of the PIC. However, when the PIC is in 10-Gigabit Ethernet mode, you must execute the **show interfaces diagnostics optics et-fpcl/picl/port:channel** command. The output of this command displays the diagnostic optics information about the corresponding 10-Gigabit Ethernet port of the PIC. For information about the P2-10G-40G-QSFPP PIC, see "[P2-10G-40G-QSFPP PIC Overview](#)" on page 182.

## Required Privilege Level

view

## Output Fields

[Table 145 on page 1370](#) lists the output fields for the **show interfaces diagnostics optics** command for DWDM and DWDM OTN PICs. Output fields are listed in the approximate order in which they appear.

**Table 145: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs**

| Field Name                | Field Description                                                                                                                                          |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b> | Name of the physical interface.                                                                                                                            |
| <b>Laser bias current</b> | Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| <b>Laser output power</b> | Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the <b>LsPOWMON</b> pin in hardware.     |

Table 145: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs (Continued)

| Field Name                            | Field Description                                                                                                                                                        |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Receiver signal average optical power | Average received optical power, in mW and dBm. This indicator is a software equivalent to the <b>RxPOWMON</b> pin in hardware. Average optical power is vendor-specific. |
| Laser end-of-life alarm               | Laser end-of-life alarm: <b>On</b> or <b>Off</b> .                                                                                                                       |
| Laser wavelength alarm                | Laser wavelength alarm: <b>On</b> or <b>Off</b> .                                                                                                                        |
| Laser bias current alarm              | Laser bias current alarm: <b>On</b> or <b>Off</b> .                                                                                                                      |
| Laser temperature alarm               | Laser temperature alarm: <b>On</b> or <b>Off</b> .                                                                                                                       |
| Laser power alarm                     | Laser power alarm: <b>On</b> or <b>Off</b> .                                                                                                                             |
| Modulator temperature alarm           | Modulator temperature alarm: <b>On</b> or <b>Off</b> .<br>Transceivers from some vendors do not support this field.                                                      |
| Modulator bias alarm                  | Modulator bias alarm: <b>On</b> or <b>Off</b> .                                                                                                                          |
| Tx multiplexer FIFO error alarm       | Transmit multiplexer first in, first out (FIFO) error alarm: <b>On</b> or <b>Off</b> .                                                                                   |
| Tx loss of PLL lock alarm             | Transmit loss of phase-locked loop (PLL) lock alarm: <b>On</b> or <b>Off</b> .                                                                                           |

**Table 145: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs (Continued)**

| Field Name                                    | Field Description                                                                                                      |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Rx loss of average optical power alarm</b> | Receive loss of average optical power alarm: <b>On</b> or <b>Off</b> .                                                 |
| <b>Rx loss of AC power alarm</b>              | Receive loss of AC power alarm: <b>On</b> or <b>Off</b> .<br>Transceivers from some vendors do not support this field. |
| <b>Rx loss of PLL lock alarm</b>              | Receive loss of phase-locked loop (PLL) lock alarm: <b>On</b> or <b>Off</b> .                                          |

[Table 146 on page 1372](#) lists the output fields for the **show interfaces diagnostics optics** command when the router is operating with bidirectional SFP optics. Output fields are listed in the approximate order in which they appear.

**Table 146: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics**

| Field Name                | Field Description                                                                                                                                          |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b> | Name of the physical interface.                                                                                                                            |
| <b>Laser bias current</b> | Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| <b>Laser output power</b> | Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).                                                                           |
| <b>Module temperature</b> | Temperature of the optics module, in Celsius and Fahrenheit.                                                                                               |
| <b>Module voltage</b>     | Internally measured module voltage.                                                                                                                        |

**Table 146: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (Continued)**

| Field Name                                   | Field Description                                                                |
|----------------------------------------------|----------------------------------------------------------------------------------|
| <b>Receiver signal average optical power</b> | Average received optical power, in mW and dBm.                                   |
| <b>Wavelength Channel number</b>             | Wavelength channel number set in the optics module.                              |
| <b>Wavelength setpoint</b>                   | Wavelength set in the optics module.                                             |
| <b>Tx Dither</b>                             | Transmit dither status. Displays whether transmit dither is enabled or disabled. |
| <b>Frequency Error</b>                       | Frequency error reported from optics module.                                     |
| <b>Wavelength Error</b>                      | Wavelength error reported from optics module.                                    |
| <b>Laser bias current high alarm</b>         | Laser bias power setting high alarm. Displays <b>on</b> or <b>off</b> .          |
| <b>Laser bias current low alarm</b>          | Laser bias power setting low alarm. Displays <b>on</b> or <b>off</b> .           |
| <b>Laser bias current high warning</b>       | Laser bias power setting high warning. Displays <b>on</b> or <b>off</b> .        |
| <b>Laser bias current low warning</b>        | Laser bias power setting low warning. Displays <b>on</b> or <b>off</b> .         |
| <b>Laser output power high alarm</b>         | Laser output power high alarm. Displays <b>on</b> or <b>off</b> .                |

**Table 146: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (Continued)**

| Field Name                             | Field Description                                                   |
|----------------------------------------|---------------------------------------------------------------------|
| <b>Laser output power low alarm</b>    | Laser output power low alarm. Displays <b>on</b> or <b>off</b> .    |
| <b>Laser output power high warning</b> | Laser output power high warning. Displays <b>on</b> or <b>off</b> . |
| <b>Laser output power low warning</b>  | Laser output power low warning. Displays <b>on</b> or <b>off</b> .  |
| <b>Module temperature high alarm</b>   | Module temperature high alarm. Displays <b>on</b> or <b>off</b> .   |
| <b>Module temperature low alarm</b>    | Module temperature low alarm. Displays <b>on</b> or <b>off</b> .    |
| <b>Module temperature high warning</b> | Module temperature high warning. Displays <b>on</b> or <b>off</b> . |
| <b>Module temperature low warning</b>  | Module temperature low warning. Displays <b>on</b> or <b>off</b> .  |
| <b>Module voltage high alarm</b>       | Module voltage high alarm. Displays <b>on</b> or <b>off</b> .       |

**Table 146: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (Continued)**

| Field Name                         | Field Description                                                                                |
|------------------------------------|--------------------------------------------------------------------------------------------------|
| <b>Module voltage low alarm</b>    | Module voltage low alarm. Displays <b>on</b> or <b>off</b> .                                     |
| <b>Module voltage high warning</b> | Module voltage high warning. Displays <b>on</b> or <b>off</b> .                                  |
| <b>Module voltage low warning</b>  | Module voltage high warning. Displays <b>on</b> or <b>off</b> .                                  |
| <b>Laser rx power high alarm</b>   | Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .                               |
| <b>Laser rx power low alarm</b>    | Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .                                |
| <b>Laser rx power high warning</b> | Receive laser power high warning. Displays <b>on</b> or <b>off</b> .                             |
| <b>Laser rx power low warning</b>  | Receive laser power low warning. Displays <b>on</b> or <b>off</b> .                              |
| <b>TEC fault alarm</b>             | TEC fault alarm. Displays <b>on</b> or <b>off</b> .                                              |
| <b>Wavelength unlocked alarm</b>   | Wavelength unlocked alarm. Displays <b>on</b> or <b>off</b> .                                    |
| <b>TxTune</b>                      | Optical transmit side status. Displays whether optical transmit side is not ready due to tuning. |

**Table 146: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (Continued)**

| Field Name                                       | Field Description                                                                                          |
|--------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Laser bias current high alarm threshold</b>   | Vendor-specified threshold for the laser bias current high alarm: <b>70.000 mA</b> .                       |
| <b>Laser bias current low alarm threshold</b>    | Vendor-specified threshold for the laser bias current low alarm: <b>0.0002 mA</b> .                        |
| <b>Laser bias current high warning threshold</b> | Vendor-specified threshold for the laser bias current high warning: <b>65.000 mA</b> .                     |
| <b>Laser bias current low warning threshold</b>  | Vendor-specified threshold for the laser bias current low warning: <b>0.0002 mA</b> .                      |
| <b>Laser output power high alarm threshold</b>   | Vendor-specified threshold for the laser output power high alarm: <b>1.0000 mW</b> or <b>0.00 dBm</b> .    |
| <b>Laser output power low alarm threshold</b>    | Vendor-specified threshold for the laser output power low alarm: <b>0.0560 mW</b> or <b>-12.52 dBm</b> .   |
| <b>Laser output power high warning threshold</b> | Vendor-specified threshold for the laser output power high warning: <b>0.6300 mW</b> or <b>-2.01 dBm</b> . |
| <b>Laser output power low warning threshold</b>  | Vendor-specified threshold for the laser output power low warning: <b>0.0890 mW</b> or <b>-10.51 dBm</b> . |

**Table 146: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (Continued)**

| Field Name                                       | Field Description                                                                             |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------|
| <b>Module temperature high alarm threshold</b>   | Vendor-specified threshold for the module temperature high alarm: <b>100° C or 212° F.</b>    |
| <b>Module temperature low alarm threshold</b>    | Vendor-specified threshold for the module temperature low alarm: <b>-50° C or -58° F.</b>     |
| <b>Module temperature high warning threshold</b> | Vendor-specified threshold for the module temperature high warning: <b>95 ° C or 203 ° F.</b> |
| <b>Module temperature low warning threshold</b>  | Vendor-specified threshold for the module temperature low warning: <b>-48° C or -54° F.</b>   |
| <b>Module voltage high alarm threshold</b>       | Module voltage high alarm threshold: <b>3.700 v.</b>                                          |
| <b>Module voltage low alarm threshold</b>        | Module voltage low alarm threshold: <b>2.900 v.</b>                                           |
| <b>Module voltage high warning threshold</b>     | Module voltage high warning threshold: <b>3.7600 v.</b>                                       |
| <b>Module voltage low warning threshold</b>      | Module voltage low warning threshold: <b>3.000 v.</b>                                         |

**Table 146: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (Continued)**

| Field Name                                   | Field Description                                                                                      |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------|
| <b>Laser rx power high alarm threshold</b>   | Vendor-specified threshold for the laser Rx power high alarm: <b>1.9953 mW</b> or <b>3.00 dBm</b> .    |
| <b>Laser rx power low alarm threshold</b>    | Vendor-specified threshold for the laser Rx power low alarm: <b>0.0001 mW</b> or <b>-40.00 dBm</b> .   |
| <b>Laser rx power high warning threshold</b> | Vendor-specified threshold for the laser Rx power high warning: <b>1.0000 mW</b> or <b>0.00 dBm</b> .  |
| <b>Laser rx power low warning threshold</b>  | Vendor-specified threshold for the laser Rx power low warning: <b>0.0010 mW</b> or <b>-30.00 dBm</b> . |

[Table 147 on page 1378](#) lists the output fields for the **show interfaces diagnostics optics** command for SFP transceivers. Output fields are listed in the approximate order in which they appear.

**Table 147: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers**

| Field Name                | Field Description                       |
|---------------------------|-----------------------------------------|
| <b>Physical interface</b> | Name of the physical interface.         |
| <b>Laser bias current</b> | Measured laser bias current in uA.      |
| <b>Laser output power</b> | Measured laser output power in mW.      |
| <b>Module temperature</b> | Internally measured module temperature. |

Table 147: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (Continued)

| Field Name                             | Field Description                                                                              |
|----------------------------------------|------------------------------------------------------------------------------------------------|
| <b>Module voltage</b>                  | Internally measured module voltage.                                                            |
| <b>Laser rx power</b>                  | Measured receive optical power in mW.                                                          |
| <b>Laser bias current high alarm</b>   | Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Laser bias current low alarm</b>    | Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| <b>Laser output power high alarm</b>   | Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Laser output power low alarm</b>    | Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| <b>Module temp high alarm</b>          | Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Module temp low alarm</b>           | Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| <b>Laser rx power high alarm</b>       | Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.    |
| <b>Laser rx power low alarm</b>        | Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Laser bias current high warning</b> | Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |

Table 147: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (Continued)

| Field Name                              | Field Description                                                                               |
|-----------------------------------------|-------------------------------------------------------------------------------------------------|
| Laser bias current low warning          | Laser bias current low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.   |
| Laser output power high warning         | Laser output power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| Laser output power low warning          | Laser output power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.   |
| Module temperature high warning         | Module temperature high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| Module temperature low warning          | Module temperature low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.   |
| Laser rx power high warning             | Laser receive power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| Laser rx power low warning              | Laser receive power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| Laser bias current high alarm threshold | Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Laser bias current low alarm threshold  | Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.             |

Table 147: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(Continued)*

| Field Name                                       | Field Description                                                                    |
|--------------------------------------------------|--------------------------------------------------------------------------------------|
| <b>Laser bias current high warning threshold</b> | Laser bias current high warning threshold. Warning ranges are vendor-specific.       |
| <b>Laser bias current low warning threshold</b>  | Laser bias current low warning threshold. Warning ranges are vendor-specific.        |
| <b>Laser output power high alarm threshold</b>   | Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific. |
| <b>Laser output power low alarm threshold</b>    | Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.  |
| <b>Laser output power high warning threshold</b> | Laser output power high warning threshold. Warning ranges are vendor-specific.       |
| <b>Laser output power low warning threshold</b>  | Laser output power low warning threshold. Warning ranges are vendor-specific.        |
| <b>Module temperature high alarm threshold</b>   | Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific. |
| <b>Module temperature low alarm threshold</b>    | Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.  |

Table 147: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers *(Continued)*

| Field Name                                       | Field Description                                                                     |
|--------------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Module temperature high warning threshold</b> | Module temperature high warning threshold. Warning ranges are vendor-specific.        |
| <b>Module temperature low warning threshold</b>  | Module temperature low warning threshold. Warning ranges are vendor-specific.         |
| <b>Module voltage high alarm threshold</b>       | Module voltage high alarm threshold. Alarm ranges are vendor-specific.                |
| <b>Module voltage low alarm threshold</b>        | Module voltage low alarm threshold. Alarm ranges are vendor-specific.                 |
| <b>Module voltage high warning threshold</b>     | Module voltage high warning threshold. Warning ranges are vendor-specific.            |
| <b>Module voltage low warning threshold</b>      | Module voltage low warning threshold. Warning ranges are vendor-specific.             |
| <b>Laser rx power high alarm threshold</b>       | Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific. |
| <b>Laser rx power low alarm threshold</b>        | Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.  |

Table 147: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (Continued)

| Field Name                                   | Field Description                                                                         |
|----------------------------------------------|-------------------------------------------------------------------------------------------|
| <b>Laser rx power high warning threshold</b> | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific. |
| <b>Laser rx power high low threshold</b>     | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific. |

Table 148 on page 1383 lists the output fields for the **show interfaces diagnostics optics** command for 10-Gigabit Ethernet transceivers. Output fields are listed in the approximate order in which they appear.

Table 148: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers

| Field Name                           | Field Description                                                                          |
|--------------------------------------|--------------------------------------------------------------------------------------------|
| <b>Physical interface</b>            | Name of the physical interface.                                                            |
| <b>Laser bias current</b>            | Measured laser bias current in mA.                                                         |
| <b>Laser output power</b>            | Measured laser output power in mW.                                                         |
| <b>Module temperature</b>            | Internally measured module temperature.                                                    |
| <b>Laser rx power</b>                | Measured receive optical power in mW.                                                      |
| <b>Laser bias current high alarm</b> | Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific. |
| <b>Laser bias current low alarm</b>  | Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.  |

**Table 148: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers**  
*(Continued)*

| Field Name                             | Field Description                                                                              |
|----------------------------------------|------------------------------------------------------------------------------------------------|
| <b>Laser output power high alarm</b>   | Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Laser output power low alarm</b>    | Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| <b>Module temp high alarm</b>          | Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Module temp low alarm</b>           | Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| <b>Laser rx power high alarm</b>       | Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.    |
| <b>Laser rx power low alarm</b>        | Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Laser bias current high warning</b> | Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| <b>Laser bias current low warning</b>  | Laser bias current low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| <b>Laser output power high warning</b> | Laser output power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| <b>Laser output power low warning</b>  | Laser output power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |

**Table 148: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers**  
*(Continued)*

| Field Name                                     | Field Description                                                                               |
|------------------------------------------------|-------------------------------------------------------------------------------------------------|
| <b>Module temperature high warning</b>         | Module temperature high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| <b>Module temperature low warning</b>          | Module temperature low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.   |
| <b>Laser rx power high warning</b>             | Laser receive power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| <b>Laser rx power low warning</b>              | Laser receive power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| <b>Laser bias current high alarm threshold</b> | Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| <b>Laser bias current low alarm threshold</b>  | Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.             |
| <b>Laser output power high alarm threshold</b> | Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| <b>Laser output power low alarm threshold</b>  | Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.             |

**Table 148: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers**  
*(Continued)*

| Field Name                                       | Field Description                                                                     |
|--------------------------------------------------|---------------------------------------------------------------------------------------|
| <b>Module temperature high alarm threshold</b>   | Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.  |
| <b>Module temperature low alarm threshold</b>    | Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.   |
| <b>Laser rx power high alarm threshold</b>       | Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific. |
| <b>Laser rx power low alarm threshold</b>        | Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.  |
| <b>Laser bias current high warning threshold</b> | Laser bias current high warning threshold. Warning ranges are vendor-specific.        |
| <b>Laser bias current low warning threshold</b>  | Laser bias current low warning threshold. Warning ranges are vendor-specific.         |
| <b>Laser output power high warning threshold</b> | Laser output power high warning threshold. Warning ranges are vendor-specific.        |
| <b>Laser output power low warning threshold</b>  | Laser output power low warning threshold. Warning ranges are vendor-specific.         |

**Table 148: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers**  
(Continued)

| Field Name                                       | Field Description                                                                         |
|--------------------------------------------------|-------------------------------------------------------------------------------------------|
| <b>Module temperature high warning threshold</b> | Module temperature high warning threshold. Warning ranges are vendor-specific.            |
| <b>Module temperature low warning threshold</b>  | Module temperature low warning threshold. Warning ranges are vendor-specific.             |
| <b>Laser rx power high warning threshold</b>     | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific. |
| <b>Laser rx power low warning threshold</b>      | Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.  |

[Table 149 on page 1387](#) lists the output fields for the **show interfaces diagnostics optics** command for 100-Gigabit Ethernet transceivers. Output fields are listed in the approximate order in which they appear.

**Table 149: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers**

| Field Name                   | Field Description                                                                                                                                                                             |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>    | Name of the physical interface.                                                                                                                                                               |
| <b>Grid Channel Number</b>   | Set of wavelengths are divided into 40 grids. At a time transceiver will be programmed in one of this set (grid). Mapping of center wavelength to grid number is presented by this parameter. |
| <b>Corrected Error Ratio</b> | Indicates accumulated Bit Error Ratio.                                                                                                                                                        |

**Table 149: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers**  
*(Continued)*

| Field Name                           | Field Description                                                                                                   |
|--------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Uncorrected Words Ratio              | Monitors the error rate for either the full 100G link or on a channel-by-channel basis. Indicates the frame errors. |
| <b>Laser bias current</b>            | Measured laser bias current in mA.                                                                                  |
| <b>Laser output power</b>            | Measured laser output power in mW.                                                                                  |
| <b>Module temperature</b>            | Internally measured module temperature.                                                                             |
| <b>Laser rx power</b>                | Measured receive optical power in mW.                                                                               |
| <b>Laser bias current high alarm</b> | Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.                          |
| <b>Laser bias current low alarm</b>  | Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.                           |
| <b>Laser output power high alarm</b> | Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.                          |
| <b>Laser output power low alarm</b>  | Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.                           |
| <b>Module temp high alarm</b>        | Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.                          |
| <b>Module temp low alarm</b>         | Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.                           |

**Table 149: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers**  
*(Continued)*

| Field Name                             | Field Description                                                                               |
|----------------------------------------|-------------------------------------------------------------------------------------------------|
| <b>Laser rx power high alarm</b>       | Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| <b>Laser rx power low alarm</b>        | Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| <b>Laser bias current high warning</b> | Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| <b>Laser bias current low warning</b>  | Laser bias current low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.   |
| <b>Laser output power high warning</b> | Laser output power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| <b>Laser output power low warning</b>  | Laser output power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.   |
| <b>Module temperature high warning</b> | Module temperature high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.  |
| <b>Module temperature low warning</b>  | Module temperature low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific.   |
| <b>Laser rx power high warning</b>     | Laser receive power high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |

**Table 149: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers**  
*(Continued)*

| Field Name                                     | Field Description                                                                              |
|------------------------------------------------|------------------------------------------------------------------------------------------------|
| <b>Laser rx power low warning</b>              | Laser receive power low warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| <b>Laser bias current high alarm threshold</b> | Laser bias current high alarm threshold. Alarm threshold ranges are vendor-specific.           |
| <b>Laser bias current low alarm threshold</b>  | Laser bias current low alarm threshold. Alarm threshold ranges are vendor-specific.            |
| <b>Laser output power high alarm threshold</b> | Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.           |
| <b>Laser output power low alarm threshold</b>  | Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.            |
| <b>Module temperature high alarm threshold</b> | Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.           |
| <b>Module temperature low alarm threshold</b>  | Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.            |
| <b>Laser rx power high alarm threshold</b>     | Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.          |

**Table 149: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers**  
*(Continued)*

| Field Name                                       | Field Description                                                                         |
|--------------------------------------------------|-------------------------------------------------------------------------------------------|
| <b>Laser rx power low alarm threshold</b>        | Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.      |
| <b>Laser bias current high warning threshold</b> | Laser bias current high warning threshold. Warning ranges are vendor-specific.            |
| <b>Laser bias current low warning threshold</b>  | Laser bias current low warning threshold. Warning ranges are vendor-specific.             |
| <b>Laser output power high warning threshold</b> | Laser output power high warning threshold. Warning ranges are vendor-specific.            |
| <b>Laser output power low warning threshold</b>  | Laser output power low warning threshold. Warning ranges are vendor-specific.             |
| <b>Module temperature high warning threshold</b> | Module temperature high warning threshold. Warning ranges are vendor-specific.            |
| <b>Module temperature low warning threshold</b>  | Module temperature low warning threshold. Warning ranges are vendor-specific.             |
| <b>Laser rx power high warning threshold</b>     | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific. |

**Table 149: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers**  
(Continued)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Laser rx power low warning threshold</b> | Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.                                                                                                                                                                                                                                                                                                                                            |
| Lane carrier frequency offset               | Difference (in frequency units) between the target frequency and the actual frequency.                                                                                                                                                                                                                                                                                                                                              |
| Lane SNR                                    | Signal-to-noise ratio of the electrical data present on the channel.                                                                                                                                                                                                                                                                                                                                                                |
| TEC Current                                 | Monitors the amount of current flowing to the TEC of a cooled laser. It is a 16-bit signed 2s complement value with a LSB unit of 0.1 mA. Thus the total range is from -3.2768 A to +3.2767 A.                                                                                                                                                                                                                                      |
| Residual ISI                                | Measures the amount of correction being done by the module to account for residual inter- symbol interference (ISI). The usual cause for this is optical dispersion so this measurement is a proxy for residual (uncorrected) optical dispersion that is being corrected by the module. The parameter is unitless and the threshold alarm and warning values will give an indication of the severity of the uncorrected dispersion. |
| PAM Histogram                               | Provides the rate of measured signal on the line that has an analog level near the cutoff for a PAM bit transition (for example, 0 <-> 1, 1<->2, 2<->3).                                                                                                                                                                                                                                                                            |

[Table 150 on page 1392](#) lists the output fields for the **show interfaces diagnostics optics** command for XFP transceivers. Output fields are listed in the approximate order in which they appear.

**Table 150: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers**

| Field Name                | Field Description               |
|---------------------------|---------------------------------|
| <b>Physical interface</b> | Name of the physical interface. |

**Table 150: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers**  
*(Continued)*

| Field Name                             | Field Description                                                                                                                                          |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Laser bias current</b>              | Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| <b>Laser output power</b>              | Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm). This is a software equivalent to the <b>LsPOWMON</b> pin in hardware.     |
| <b>Module temperature</b>              | Temperature of the XFP optics module, in Celsius and Fahrenheit.                                                                                           |
| <b>Laser rx power</b>                  | Laser received optical power, in mW and dBm.                                                                                                               |
| <b>Laser bias current high alarm</b>   | Laser bias power setting high alarm. Displays <b>on</b> or <b>off</b> .                                                                                    |
| <b>Laser bias current low alarm</b>    | Laser bias power setting low alarm. Displays <b>on</b> or <b>off</b> .                                                                                     |
| <b>Laser bias current high warning</b> | Laser bias power setting high warning. Displays <b>on</b> or <b>off</b> .                                                                                  |
| <b>Laser bias current low warning</b>  | Laser bias power setting low warning. Displays <b>on</b> or <b>off</b> .                                                                                   |
| <b>Laser output power high alarm</b>   | Laser output power high alarm. Displays <b>on</b> or <b>off</b> .                                                                                          |
| <b>Laser output power low alarm</b>    | Laser output power low alarm. Displays <b>on</b> or <b>off</b> .                                                                                           |

**Table 150: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers**  
*(Continued)*

| Field Name                             | Field Description                                                   |
|----------------------------------------|---------------------------------------------------------------------|
| <b>Laser output power high warning</b> | Laser output power high warning. Displays <b>on</b> or <b>off</b> . |
| <b>Laser output power low warning</b>  | Laser output power low warning. Displays <b>on</b> or <b>off</b> .  |
| <b>Module temperature high alarm</b>   | Module temperature high alarm. Displays <b>on</b> or <b>off</b> .   |
| <b>Module temperature low alarm</b>    | Module temperature low alarm. Displays <b>on</b> or <b>off</b> .    |
| <b>Module temperature high warning</b> | Module temperature high warning. Displays <b>on</b> or <b>off</b> . |
| <b>Module temperature low warning</b>  | Module temperature low warning. Displays <b>on</b> or <b>off</b> .  |
| <b>Laser rx power high alarm</b>       | Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .  |
| <b>Laser rx power low alarm</b>        | Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .   |

**Table 150: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers**  
(Continued)

| Field Name                         | Field Description                                                                                                                           |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Laser rx power high warning</b> | Receive laser power high warning. Displays <b>on</b> or <b>off</b> .                                                                        |
| <b>Laser rx power low warning</b>  | Receive laser power low warning. Displays <b>on</b> or <b>off</b> .                                                                         |
| <b>Module not ready alarm</b>      | Module not ready alarm. When <b>on</b> , indicates the module has an operational fault. Displays <b>on</b> or <b>off</b> .                  |
| <b>Module power down alarm</b>     | Module power down alarm. When <b>on</b> , module is in a limited power mode, low for normal operation. Displays <b>on</b> or <b>off</b> .   |
| <b>Tx data not ready alarm</b>     | Any condition leading to invalid data on the transmit path. Displays <b>on</b> or <b>off</b> .                                              |
| <b>Tx not ready alarm</b>          | Any condition leading to invalid data on the transmit path. Displays <b>on</b> or <b>off</b> .                                              |
| <b>Tx laser fault alarm</b>        | Laser fault condition. Displays <b>on</b> or <b>off</b> .                                                                                   |
| <b>Tx CDR loss of lock alarm</b>   | Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays <b>on</b> or <b>off</b> .       |
| <b>Rx not ready alarm</b>          | Any condition leading to invalid data on the receive path. Displays <b>on</b> or <b>off</b> .                                               |
| <b>Rx loss of signal alarm</b>     | Receive Loss of Signal alarm. When <b>on</b> , indicates insufficient optical input power to the module. Displays <b>on</b> or <b>off</b> . |

**Table 150: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers (Continued)**

| Field Name                                       | Field Description                                                                                         |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Rx CDR loss of lock alarm</b>                 | Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays <b>on</b> or <b>off</b> . |
| <b>Laser bias current high alarm threshold</b>   | Vendor-specified threshold for the laser bias current high alarm: <b>130.000 mA</b> .                     |
| <b>Laser bias current low alarm threshold</b>    | Vendor-specified threshold for the laser bias current low alarm: <b>10.000 mA</b> .                       |
| <b>Laser bias current high warning threshold</b> | Vendor-specified threshold for the laser bias current high warning: <b>120.000 mA</b> .                   |
| <b>Laser bias current low warning threshold</b>  | Vendor-specified threshold for the laser bias current low warning: <b>12.000 mA</b> .                     |
| <b>Laser output power high alarm threshold</b>   | Vendor-specified threshold for the laser output power high alarm: <b>0.8910 mW</b> or <b>-0.50 dBm</b> .  |
| <b>Laser output power low alarm threshold</b>    | Vendor-specified threshold for the laser output power low alarm: <b>0.2230 mW</b> or <b>-6.52 dBm</b> .   |
| <b>Laser output power high warning threshold</b> | Vendor-specified threshold for the laser output power high warning: <b>0.7940 mW</b> or <b>-100 dBm</b> . |

**Table 150: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers**  
*(Continued)*

| Field Name                                       | Field Description                                                                                        |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Laser output power low warning threshold</b>  | Vendor-specified threshold for the laser output power low warning: <b>0.2510 mW</b> or <b>-600 dBm</b> . |
| <b>Module temperature high alarm threshold</b>   | Vendor-specified threshold for the module temperature high alarm: <b>90° C</b> or <b>194° F</b> .        |
| <b>Module temperature low alarm threshold</b>    | Vendor-specified threshold for the module temperature low alarm: <b>-5° C</b> or <b>23° F</b> .          |
| <b>Module temperature high warning threshold</b> | Vendor-specified threshold for the module temperature high warning: <b>85° C</b> or <b>185° F</b> .      |
| <b>Module temperature low warning threshold</b>  | Vendor-specified threshold for the module temperature low warning: <b>0° C</b> or <b>32° F</b> .         |
| <b>Laser rx power high alarm threshold</b>       | Vendor-specified threshold for the laser Rx power high alarm: <b>1.2589 mW</b> or <b>1.00 dBm</b> .      |
| <b>Laser rx power low alarm threshold</b>        | Vendor-specified threshold for the laser Rx power low alarm: <b>0.0323 mW</b> or <b>-14.91 dBm</b> .     |
| <b>Laser rx power high warning threshold</b>     | Vendor-specified threshold for the laser Rx power high warning: <b>1.1220 mW</b> or <b>0.50 dBm</b> .    |

**Table 150: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers (Continued)**

| Field Name                                  | Field Description                                                                                      |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------|
| <b>Laser rx power low warning threshold</b> | Vendor-specified threshold for the laser Rx power low warning: <b>0.0363 mW</b> or <b>-14.40 dBm</b> . |

[Table 151 on page 1398](#) lists the output fields for the **show interfaces diagnostics optics** command for VCP. Output fields are listed in the approximate order in which they appear.

**Table 151: show interfaces diagnostics optics Output for Virtual Chassis Ports**

| Field Name                                   | Field Description                                                                                                                                          |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>                    | Name of the physical interface.                                                                                                                            |
| <b>Laser bias current</b>                    | Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| <b>Laser output power</b>                    | Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).                                                                           |
| <b>Module temperature</b>                    | Temperature of the optics module, in Celsius and Fahrenheit.                                                                                               |
| <b>Module voltage</b>                        | Internally measured module voltage.                                                                                                                        |
| <b>Receiver signal average optical power</b> | Average received optical power, in mW and dBm.                                                                                                             |
| <b>Laser bias current high alarm</b>         | Laser bias power setting high alarm. Displays <b>on</b> or <b>off</b> .                                                                                    |

Table 151: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

| Field Name                             | Field Description                                                         |
|----------------------------------------|---------------------------------------------------------------------------|
| <b>Laser bias current low alarm</b>    | Laser bias power setting low alarm. Displays <b>on</b> or <b>off</b> .    |
| <b>Laser bias current high warning</b> | Laser bias power setting high warning. Displays <b>on</b> or <b>off</b> . |
| <b>Laser bias current low warning</b>  | Laser bias power setting low warning. Displays <b>on</b> or <b>off</b> .  |
| <b>Laser output power high alarm</b>   | Laser output power high alarm. Displays <b>on</b> or <b>off</b> .         |
| <b>Laser output power low alarm</b>    | Laser output power low alarm. Displays <b>on</b> or <b>off</b> .          |
| <b>Laser output power high warning</b> | Laser output power high warning. Displays <b>on</b> or <b>off</b> .       |
| <b>Laser output power low warning</b>  | Laser output power low warning. Displays <b>on</b> or <b>off</b> .        |
| <b>Module temperature high alarm</b>   | Module temperature high alarm. Displays <b>on</b> or <b>off</b> .         |
| <b>Module temperature low alarm</b>    | Module temperature low alarm. Displays <b>on</b> or <b>off</b> .          |

Table 151: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

| Field Name                             | Field Description                                                    |
|----------------------------------------|----------------------------------------------------------------------|
| <b>Module temperature high warning</b> | Module temperature high warning. Displays <b>on</b> or <b>off</b> .  |
| <b>Module temperature low warning</b>  | Module temperature low warning. Displays <b>on</b> or <b>off</b> .   |
| <b>Module voltage high alarm</b>       | Module voltage high alarm. Displays <b>on</b> or <b>off</b> .        |
| <b>Module voltage low alarm</b>        | Module voltage low alarm. Displays <b>on</b> or <b>off</b> .         |
| <b>Module voltage high warning</b>     | Module voltage high warning. Displays <b>on</b> or <b>off</b> .      |
| <b>Module voltage low warning</b>      | Module voltage low warning. Displays <b>on</b> or <b>off</b> .       |
| <b>Laser rx power high alarm</b>       | Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .   |
| <b>Laser rx power low alarm</b>        | Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .    |
| <b>Laser rx power high warning</b>     | Receive laser power high warning. Displays <b>on</b> or <b>off</b> . |
| <b>Laser rx power low warning</b>      | Receive laser power low warning. Displays <b>on</b> or <b>off</b> .  |

Table 151: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

| Field Name                                       | Field Description                                                   |
|--------------------------------------------------|---------------------------------------------------------------------|
| <b>Laser bias current high alarm threshold</b>   | Vendor-specified threshold for the laser bias current high alarm.   |
| <b>Laser bias current low alarm threshold</b>    | Vendor-specified threshold for the laser bias current low alarm.    |
| <b>Laser bias current high warning threshold</b> | Vendor-specified threshold for the laser bias current high warning. |
| <b>Laser bias current low warning threshold</b>  | Vendor-specified threshold for the laser bias current low warning.  |
| <b>Laser output power high alarm threshold</b>   | Vendor-specified threshold for the laser output power high alarm.   |
| <b>Laser output power low alarm threshold</b>    | Vendor-specified threshold for the laser output power low alarm.    |
| <b>Laser output power high warning threshold</b> | Vendor-specified threshold for the laser output power high warning. |
| <b>Laser output power low warning threshold</b>  | Vendor-specified threshold for the laser output power low warning.  |

Table 151: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

| Field Name                                       | Field Description                                                   |
|--------------------------------------------------|---------------------------------------------------------------------|
| <b>Module temperature high alarm threshold</b>   | Vendor-specified threshold for the module temperature high alarm.   |
| <b>Module temperature low alarm threshold</b>    | Vendor-specified threshold for the module temperature low alarm.    |
| <b>Module temperature high warning threshold</b> | Vendor-specified threshold for the module temperature high warning. |
| <b>Module temperature low warning threshold</b>  | Vendor-specified threshold for the module temperature low warning.  |
| <b>Module voltage high alarm threshold</b>       | Module voltage high alarm threshold.                                |
| <b>Module voltage low alarm threshold</b>        | Module voltage low alarm threshold.                                 |
| <b>Module voltage high warning threshold</b>     | Module voltage high warning threshold.                              |
| <b>Module voltage low warning threshold</b>      | Module voltage low warning threshold.                               |

Table 151: show interfaces diagnostics optics Output for Virtual Chassis Ports *(Continued)*

| Field Name                                   | Field Description                                               |
|----------------------------------------------|-----------------------------------------------------------------|
| <b>Laser rx power high alarm threshold</b>   | Vendor-specified threshold for the laser Rx power high alarm.   |
| <b>Laser rx power low alarm threshold</b>    | Vendor-specified threshold for the laser Rx power low alarm.    |
| <b>Laser rx power high warning threshold</b> | Vendor-specified threshold for the laser Rx power high warning. |
| <b>Laser rx power low warning threshold</b>  | Vendor-specified threshold for the laser Rx power low warning.  |

## Sample Output

**show interfaces diagnostics optics for 10-Gigabit Ethernet on JNP10K-LC1201 line card (PTX10008)**

```

user@host> show interfaces diagnostics optics et-4/0/20
Physical interface: et-4/0/20
 Module temperature : 34 degrees C / 93 degrees F
 Module voltage : 3.315 V
 Wavelength channel number : 42
 Wavelength setpoint : 1550.10 nm
 Tx dither : Enabled
 Frequency error : 0.20 GHz
 Wavelength error : 0 nm
 TEC fault alarm : False
 Wavelength unlocked alarm : False
 Tx tune alarm : False
 Module temperature high alarm : Off

```

```

Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Cont.....

```

### show interfaces diagnostics optics (DWDM and DWDM OTN)

```

user@host> show interfaces diagnostics optics ge-5/0/0
Physical interface: ge-5/0/0
Laser bias current : 79.938 mA
Laser output power : 1.592 mW / 2.02 dBm
Receiver signal average optical power : 1.3854 mW / 1.42 dBm
Laser end-of-life alarm : Off
Laser wavelength alarm : Off
Laser bias current alarm : Off
Laser temperature alarm : Off
Laser power alarm : Off
Modulator temperature alarm : Off
Modulator bias alarm : Off
Tx multiplexer FIFO error alarm : Off
Tx loss of PLL lock alarm : Off
Rx loss of average optical power alarm : Off
Rx loss of AC power alarm : Off
Rx loss of PLL lock alarm : Off

```

### show interfaces diagnostics optics (MPC6E with OTN MIC)

```

user@host> show interfaces diagnostics optics xe-3/0/0
Physical interface: xe-3/0/0
Laser bias current : 7.806 mA
Laser output power : 0.5660 mW / -2.47 dBm
Module temperature : 32 degrees C / 89 degrees F
Module voltage : 3.3560 V
Receiver signal average optical power : 0.5501 mW / -2.60 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off

```

```
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 11.800 mA
Laser bias current low alarm threshold : 4.000 mA
Laser bias current high warning threshold : 10.800 mA
Laser bias current low warning threshold : 5.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold : 0.2510 mW / -6.00 dBm
Laser output power high warning threshold : 0.6600 mW / -1.80 dBm
Laser output power low warning threshold : 0.3160 mW / -5.00 dBm
Module temperature high alarm threshold : 78 degrees C / 172 degrees F
Module temperature low alarm threshold : -13 degrees C / 9 degrees F
Module temperature high warning threshold : 73 degrees C / 163 degrees F
Module temperature low warning threshold : -8 degrees C / 18 degrees F
Module voltage high alarm threshold : 3.700 V
Module voltage low alarm threshold : 2.900 V
Module voltage high warning threshold : 3.600 V
Module voltage low warning threshold : 3.000 V
Laser rx power high alarm threshold : 1.0000 mW / 0.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm
```

**show interfaces diagnostics optics (Bidirectional SFP)**

```

user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
 Laser bias current : 13.356 mA
 Laser output power : 0.2210 mW / -6.56 dBm
 Module temperature : 36 degrees C / 96 degrees F
 Module voltage : 3.2180 V
 Receiver signal average optical power : 0.2429 mW / -6.15 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : Off
 Laser rx power high warning : Off
 Laser rx power low warning : Off
 Laser bias current high alarm threshold : 70.000 mA
 Laser bias current low alarm threshold : 0.002 mA
 Laser bias current high warning threshold : 65.000 mA
 Laser bias current low warning threshold : 0.002 mA
 Laser output power high alarm threshold : 1.0000 mW / 0.00 dBm
 Laser output power low alarm threshold : 0.0560 mW / -12.52 dBm
 Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
 Laser output power low warning threshold : 0.0890 mW / -10.51 dBm
 Module temperature high alarm threshold : 100 degrees C / 212 degrees F
 Module temperature low alarm threshold : -50 degrees C / -58 degrees F
 Module temperature high warning threshold : 95 degrees C / 203 degrees F
 Module temperature low warning threshold : -48 degrees C / -54 degrees F
 Module voltage high alarm threshold : 3.700 V

```

```

Module voltage low alarm threshold : 2.900 V
Module voltage high warning threshold : 3.600 V
Module voltage low warning threshold : 3.000 V
Laser rx power high alarm threshold : 1.9953 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0001 mW / -40.00 dBm
Laser rx power high warning threshold : 1.0000 mW / 0.00 dBm
Laser rx power low warning threshold : 0.0010 mW / -30.00 dBm

```

### show interfaces diagnostics optics (SFP)

```

user@host> show interfaces diagnostics optics ge-0/3/0
Physical interface: ge-0/3/0
 Laser bias current : 23.408 mA
 Laser output power : 1.479 mW / 1.70 dBm
 Module temperature : 37 degrees C / 99 degrees F
 Laser rx power : 0.121 mW / -9.16 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Laser rx power high warning : Off
 Laser rx power low warning : Off
 Laser bias current high alarm threshold : 31.000 mA
 Laser bias current low alarm threshold : 10.000 mA
 Laser output power high alarm threshold : 6.000 mW / 7.78 dBm
 Laser output power low alarm threshold : 0.100 mW / -10.00 dBm
 Module temperature high alarm threshold : 85 degrees C / 185 degrees F
 Module temperature low alarm threshold : 0 degrees C / 32 degrees F
 Laser rx power high alarm threshold : 1.000 mW / 0.00 dBm
 Laser rx power low alarm threshold : 0.001 mW / -30.00 dBm
 Laser bias current high warning threshold : 28.000 mA

```

```

Laser bias current low warning threshold : 11.000 mA
Laser output power high warning threshold : 5.000 mW / 6.99 dBm
Laser output power low warning threshold : 0.500 mW / -3.01 dBm
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 10 degrees C / 50 degrees F
Laser rx power high warning threshold : 0.501 mW / -3.00 dBm
Laser rx power low warning threshold : 0.001 mW / -28.86 dBm

```

### show interfaces diagnostics optics (SFP)

```

user@host> show interfaces diagnostics optics ge-1/0/0
Physical interface: ge-1/0/0
 Laser bias current : 49.010 mA
 Laser output power : 1.263 mW / 1.01 dBm
 Module temperature : 17 degrees C / 62 degrees F
 Module voltage : 4.21 V
 Laser rx power : 0.060 mW / -12.21 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high warning : Off
 Laser rx power low warning : Off
 Laser bias current high alarm threshold : 70.000 mA
 Laser bias current low alarm threshold : 20.000 mA
 Laser bias current high warning threshold : 65.000 mA
 Laser bias current low warning threshold : 25.000 mA

```

```

Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.1990 mW / -7.01 dBm
Laser output power high warning threshold : 1.2580 mW / 1.00 dBm
Laser output power low warning threshold : 0.2230 mW / -6.52 dBm
Module temperature high alarm threshold : 78 degrees C /172 degrees F
Module temperature low alarm threshold : 13 degrees C / 9 degrees F
Module temperature high warning threshold : 75 degrees C /167 degrees F
Module temperature low warning threshold : 10 degrees C / 14 degrees F
Module voltage high alarm threshold : 5.71 V
Module voltage low alarm threshold : 2.05 V
Module voltage high warning threshold : 5.20 V
Module voltage low warning threshold : 3.11 V
Laser rx power high alarm threshold : 1.7783 mW / 2.50 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.5849 mW / 2.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

### show interfaces diagnostics optics (XFP and CFP Optics)

```

user@host> show interfaces diagnostics optics xe-2/1/0
Physical interface: xe-2/1/0
Laser bias current : 52.060 mA
Laser output power : 0.5640 mW / -2.49 dBm
Module temperature : 31 degrees C / 88 degrees F
Laser rx power : 0.0844 mW / -10.74 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off

```

```

Module not ready alarm : Off
Module power down alarm : Off
Tx data not ready alarm : Off
Tx not ready alarm : Off
Tx laser fault alarm : Off
Tx CDR loss of lock alarm : Off
Rx not ready alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
Laser bias current high alarm threshold : 130.000 mA
Laser bias current low alarm threshold : 10.000 mA
Laser bias current high warning threshold : 120.000 mA
Laser bias current low warning threshold : 12.000 mA
Laser output power high alarm threshold : 0.8910 mW / -0.50 dBm
Laser output power low alarm threshold : 0.2230 mW / -6.52 dBm
Laser output power high warning threshold : 0.7940 mW / -1.00 dBm
Laser output power low warning threshold : 0.2510 mW / -6.00 dBm
Module temperature high alarm threshold : 90 degrees C / 194 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold : 1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0323 mW / -14.91 dBm
Laser rx power high warning threshold : 1.1220 mW / 0.50 dBm
Laser rx power low warning threshold : 0.0363 mW / -14.40 dBm

```

### show interfaces diagnostics optics for 10-Gigabit Ethernet (PTX 24-10GE-SFPP)

```

user@host> show interfaces diagnostics optics et-2/0/23
Physical interface: et-2/0/23
Laser bias current : 8.482 mA
Laser output power : 0.5890 mW / -2.30 dBm
Module temperature : 51 degrees C / 123 degrees F
Module voltage : 3.2970 V
Receiver signal average optical power : 0.5574 mW / -2.54 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off

```

```

Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 11.800 mA
Laser bias current low alarm threshold : 4.000 mA
Laser bias current high warning threshold : 10.800 mA
Laser bias current low warning threshold : 5.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold : 0.2510 mW / -6.00 dBm
Laser output power high warning threshold : 0.6600 mW / -1.80 dBm
Laser output power low warning threshold : 0.3160 mW / -5.00 dBm
Module temperature high alarm threshold : 93 degrees C / 199 degrees F
Module temperature low alarm threshold : -13 degrees C / 9 degrees F
Module temperature high warning threshold : 88 degrees C / 190 degrees F
Module temperature low warning threshold : -8 degrees C / 18 degrees F
Module voltage high alarm threshold : 3.700 V
Module voltage low alarm threshold : 2.900 V
Module voltage high warning threshold : 3.600 V
Module voltage low warning threshold : 3.000 V
Laser rx power high alarm threshold : 1.0000 mW / 0.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

### show interfaces diagnostics optics for 40-Gigabit Ethernet

```

user@host> show interfaces diagnostics optics et-7/1/0
Physical interface: et-7/1/0
 Module temperature : 34 degrees C / 94 degrees F

```

```

Module voltage : 3.4720 V
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Module not ready alarm : Off
Module low power alarm : Off
Module initialization incomplete alarm : Off
Module fault alarm : Off
PLD Flash initialization fault alarm : Off
Power supply fault alarm : Off
Checksum fault alarm : Off
Tx laser disabled alarm : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
Module temperature high alarm threshold : 80 degrees C / 176 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 75 degrees C / 167 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.5990 V
Module voltage low alarm threshold : 3.0000 V
Module voltage high warning threshold : 3.5000 V
Module voltage low warning threshold : 3.0990 V
Laser bias current high alarm threshold : 100.000 mA
Laser bias current low alarm threshold : 10.000 mA
Laser bias current high warning threshold : 80.000 mA
Laser bias current low warning threshold : 15.000 mA
Laser output power high alarm threshold : 2.8180 mW / 4.50 dBm
Laser output power low alarm threshold : 0.2390 mW / -6.22 dBm
Laser output power high warning threshold : 2.2380 mW / 3.50 dBm
Laser output power low warning threshold : 0.3010 mW / -5.21 dBm
Laser rx power high alarm threshold : 2.5119 mW / 4.00 dBm
Laser rx power low alarm threshold : 0.0316 mW / -15.00 dBm
Laser rx power high warning threshold : 1.9953 mW / 3.00 dBm
Laser rx power low warning threshold : 0.0631 mW / -12.00 dBm
Laser temperature high alarm threshold : 80 degrees C / 176 degrees F
Laser temperature low alarm threshold : -10 degrees C / 14 degrees F

```

```

Laser temperature high warning threshold : 75 degrees C / 167 degrees F
Laser temperature low warning threshold : -5 degrees C / 23 degrees F
Lane 0
Laser bias current : 27.829 mA
Laser output power : 0.851 mW / -0.70 dBm
Laser temperature : 34 degrees C / 94 degrees F
Laser receiver power : 0.894 mW / -0.49 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off
Lane 1
Laser bias current : 35.374 mA
Laser output power : 0.896 mW / -0.48 dBm
Laser temperature : 34 degrees C / 94 degrees F
Laser receiver power : 0.707 mW / -1.50 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off

```

```

Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

```

## Lane 2

```

Laser bias current : 29.173 mA
Laser output power : 0.890 mW / -0.51 dBm
Laser temperature : 34 degrees C / 94 degrees F
Laser receiver power : 0.704 mW / -1.52 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

```

```

Lane 3
Laser bias current : 36.164 mA
Laser output power : 0.899 mW / -0.46 dBm
Laser temperature : 34 degrees C / 94 degrees F
Laser receiver power : 0.892 mW / -0.50 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

```

### show interfaces diagnostics optics (P1-PTX-2-100G-WDM)

```

user@host> show interfaces diagnostics optics et-1/0/0
Physical interface: et-1/0/0
Module temperature : 37 degrees C / 98 degrees F
Module voltage : 3.3370 V
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off

```

```

Module voltage high warning : Off
Module voltage low warning : Off
Module not ready alarm : Off
Module low power alarm : Off
Module initialization incomplete alarm : Off
Module fault alarm : Off
PLD Flash initialization fault alarm : Off
Power supply fault alarm : Off
Checksum fault alarm : Off
Tx laser disabled alarm : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
Module temperature high alarm threshold : 70 degrees C / 158 degrees F
Module temperature low alarm threshold : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 68 degrees C / 154 degrees F
Module temperature low warning threshold : 2 degrees C / 36 degrees F
Module voltage high alarm threshold : 3.4640 V
Module voltage low alarm threshold : 3.1340 V
Module voltage high warning threshold : 3.4310 V
Module voltage low warning threshold : 3.1670 V
Laser bias current high alarm threshold : 300.000 mA
Laser bias current low alarm threshold : 75.000 mA
Laser bias current high warning threshold : 287.500 mA
Laser bias current low warning threshold : 87.500 mA
Rx power high alarm threshold : 2.8184 mW / 4.50 dBm
Rx power low alarm threshold : 0.0251 mW / -16.00 dBm
Rx power high warning threshold : 2.5119 mW / 4.00 dBm
Rx power low warning threshold : 0.0501 mW / -13.00 dBm
LOS alarm threshold : 0.0158mW/ -18.01 dBm
LOS warning threshold : 0.0251mW/ -16.00 DBm
Laser temperature high alarm threshold : 57 degrees C / 135 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 : 27 degrees C / 81 degrees F
Lane 0
Laser bias current : 164.384 mA
Tx power : 1.181 mW / 0.72 dBm
Laser temperature : 41 degrees C / 106 degrees F
Rx power : 0.632 mW / -1.99 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off

```

```

Laser bias current high warning : Off
Laser bias current low warning : Off
Tx power high alarm : Off
Tx power low alarm : Off
Tx power high warning : Off
Tx power low warning : Off
Laser temperature high alarm : Off
Laser temperature low alarm : Off
Laser temperature high warning : Off
Laser temperature low warning : Off
Rx power high alarm : Off
Rx power low alarm : Off
Rx power high warning : Off
Rx power low warning : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
APD supply fault alarm : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off

```

### show interfaces diagnostics optics (P1-PTX-24-10G-W-SFPP )

```

user@host> show interfaces diagnostics optics ge-3/0/6
Physical interface: ge-3/0/6
Laser bias current : 13.356 mA
Laser output power : 0.2210 mW / -6.56 dBm
Module temperature : 36 degrees C / 96 degrees F
Module voltage : 3.2180 V
Receiver signal average optical power : 0.2429 mW / -6.15 dBm
Wavelength Channel number : 1
Wavelength setpoint : 1568.80 nm
Tx Dither : Disabled
Frequency Error : 0.00 GHz
Wavelength Error : 0.00 nm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off

```

```

Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
TEC fault alarm : Off
Wavelength unlocked alarm : Off
Tx Tune : Off
Laser bias current high alarm threshold : 70.000 mA
Laser bias current low alarm threshold : 0.002 mA
Laser bias current high warning threshold : 65.000 mA
Laser bias current low warning threshold : 0.002 mA
Laser output power high alarm threshold : 1.0000 mW / 0.00 dBm
Laser output power low alarm threshold : 0.0560 mW / -12.52 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0890 mW / -10.51 dBm
Module temperature high alarm threshold : 100 degrees C / 212 degrees F
Module temperature low alarm threshold : -50 degrees C / -58 degrees F
Module temperature high warning threshold : 95 degrees C / 203 degrees F
Module temperature low warning threshold : -48 degrees C / -54 degrees F
Module voltage high alarm threshold : 3.700 V
Module voltage low alarm threshold : 2.900 V
Module voltage high warning threshold : 3.600 V
Module voltage low warning threshold : 3.000 V
Laser rx power high alarm threshold : 1.9953 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0001 mW / -40.00 dBm
Laser rx power high warning threshold : 1.0000 mW / 0.00 dBm
Laser rx power low warning threshold : 0.0010 mW / -30.00 dBm

```

**show interfaces diagnostics optics (P2-10G-40G-QSFPP PIC in 40-Gigabit Ethernet mode)**

```

user@host> show interfaces diagnostics optics et-0/1/5
Physical interface: et-0/1/5
 Module temperature : 30 degrees C / 85 degrees F
 Module voltage : 3.2760 V
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Module temperature high alarm threshold : 75 degrees C / 167 degrees F
 Module temperature low alarm threshold : 5 degrees C / 41 degrees F
 Module temperature high warning threshold : 70 degrees C / 158 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F
 Module voltage high alarm threshold : 3.6300 V
 Module voltage low alarm threshold : 2.9700 V
 Module voltage high warning threshold : 3.4640 V
 Module voltage low warning threshold : 3.1340 V
 Laser bias current high alarm threshold : 10.000 mA
 Laser bias current low alarm threshold : 0.500 mA
 Laser bias current high warning threshold : 9.500 mA
 Laser bias current low warning threshold : 1.000 mA
 Laser output power high alarm threshold : 0.0000 mW / - Inf dBm
 Laser output power low alarm threshold : 0.0000 mW / - Inf dBm
 Laser output power high warning threshold : 0.0000 mW / - Inf dBm
 Laser output power low warning threshold : 0.0000 mW / - Inf dBm
 Laser rx power high alarm threshold : 2.1878 mW / 3.40 dBm
 Laser rx power low alarm threshold : 0.0446 mW / -13.51 dBm
 Laser rx power high warning threshold : 1.7378 mW / 2.40 dBm
 Laser rx power low warning threshold : 0.1122 mW / -9.50 dBm
Lane 0
 Laser bias current : 7.065 mA
 Laser output power : 0.710 mW / -1.49 dBm
 Laser receiver power : 0.472 mW / -3.26 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off

```

```

Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off

```

## Lane 1

```

Laser bias current : 6.978 mA
Laser output power : 0.771 mW / -1.13 dBm
Laser receiver power : 0.450 mW / -3.47 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off

```

## Lane 2

```

Laser bias current : 6.955 mA
Laser output power : 0.760 mW / -1.19 dBm
Laser receiver power : 0.556 mW / -2.55 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off

```

## Lane 3

```

Laser bias current : 6.981 mA
Laser output power : 0.736 mW / -1.33 dBm
Laser receiver power : 0.537 mW / -2.70 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off

```

```

Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off

```

### show interfaces diagnostics optics (P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode)

```

user@host> show interfaces diagnostics optics et-0/1/5:3
Physical interface: et-0/1/5:3
 Module temperature : 30 degrees C / 85 degrees F
 Module voltage : 3.2760 V
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Module temperature high alarm threshold : 75 degrees C / 167 degrees F
 Module temperature low alarm threshold : 5 degrees C / 41 degrees F
 Module temperature high warning threshold : 70 degrees C / 158 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F
 Module voltage high alarm threshold : 3.6300 V
 Module voltage low alarm threshold : 2.9700 V
 Module voltage high warning threshold : 3.4640 V
 Module voltage low warning threshold : 3.1340 V
 Laser bias current high alarm threshold : 10.000 mA
 Laser bias current low alarm threshold : 0.500 mA
 Laser bias current high warning threshold : 9.500 mA
 Laser bias current low warning threshold : 1.000 mA
 Laser output power high alarm threshold : 0.0000 mW / - Inf dBm
 Laser output power low alarm threshold : 0.0000 mW / - Inf dBm
 Laser output power high warning threshold : 0.0000 mW / - Inf dBm
 Laser output power low warning threshold : 0.0000 mW / - Inf dBm
 Laser rx power high alarm threshold : 2.1878 mW / 3.40 dBm
 Laser rx power low alarm threshold : 0.0446 mW / -13.51 dBm
 Laser rx power high warning threshold : 1.7378 mW / 2.40 dBm
 Laser rx power low warning threshold : 0.1122 mW / -9.50 dBm

```

Lane 3

```

Laser bias current : 6.981 mA
Laser output power : 0.736 mW / -1.33 dBm
Laser receiver power : 0.537 mW / -2.70 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off

```

### show interfaces diagnostics optics (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```

user@host> show interfaces diagnostics optics et-2/0/0
Physical interface: et-2/0/0
 Module temperature : 39 degrees C / 102 degrees F
 Module voltage : 3.2300 V
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Module not ready alarm : Off
 Module low power alarm : Off
 Module initialization incomplete alarm : Off
 Module fault alarm : Off
 Tx laser disabled alarm : Off
 Rx loss of signal alarm : Off
 Module temperature high alarm threshold : 90 degrees C / 194 degrees F
 Module temperature low alarm threshold : -20 degrees C / -4 degrees F
 Module temperature high warning threshold : 0 degrees C / 32 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F
 Module voltage high alarm threshold : 3.6300 V
 Module voltage low alarm threshold : 2.9700 V

```

```

Module voltage high warning threshold : 0.0000 V
Module voltage low warning threshold : 0.0000 V
Rx power high alarm threshold : 6.5535 mW / 8.16 dBm
Rx power low alarm threshold : 0.0028 mW / -25.53 dBm
Rx power high warning threshold : 6.5535 mW / 8.16 dBm
Rx power low warning threshold : 0.0028 mW / -25.53 dBm
LOS alarm threshold : 0.0028 mW / -25.53 dBm
LOS warning threshold : 0.0028 mW / -25.53 dBm
Modem lock state : OK
Lane 0
Tx power : 1.000 mW / 0.00 dBm
Module temperature : 51 degrees C / 124 degrees F
Rx power (total) : 0.644 mW / -1.91 dBm
Rx power (signal) : 0.618 mW / -2.09 dBm
Lane chromatic dispersion : -22 ps/nm
Lane differential group delay : 5 ps
Lane Q2 factor : 14.20 dB
Lane carrier frequency offset : -534 Mz
Lane electrical SNR : 9.20 dB
Tx power high alarm : Off
Tx power low alarm : Off
Tx power high warning : Off
Tx power low warning : Off
Rx power high alarm : Off
Rx power low alarm : Off
Rx power high warning : Off
Rx power low warning : Off
Rx loss of signal alarm : Off
Wavelength unlocked alarm : Off

```

### show interfaces diagnostics optics (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

```

user@host> show interfaces diagnostics optics et-4/0/0
Physical interface: et-4/0/0
Laser output power : 54 degrees C / 129 degrees F
Tx module temperature : 0.0000
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off

```

```

Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Module not ready alarm : Off
Module low power alarm : Off
Module initialization incomplete alarm : Off
Module fault alarm : Off
PLD Flash initialization fault alarm : Off
Power supply fault alarm : Off
Checksum fault alarm : Off
Tx laser disabled alarm : Off
Rx loss of signal alarm : Off
Module temperature high alarm threshold : 80 degrees C / 176 degrees F
Module temperature low alarm threshold : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 65 degrees C / 149 degrees F
Module temperature low warning threshold : 5 degrees C / 41 degrees F
Module voltage high alarm threshold : 0.0000 V
Module voltage low alarm threshold : 0.0000 V
Module voltage high warning threshold : 0.0000 V
Module voltage low warning threshold : 0.0000 V
Rx power high alarm threshold : 0.0000 mW / - Inf dBm
Rx power low alarm threshold : 0.0000 mW / - Inf dBm
Rx power high warning threshold : 0.0000 mW / - Inf dBm
Rx power low warning threshold : 0.0000 mW / - Inf dBm
LOS alarm threshold : 0.0158 mW / -18.01 dBm
LOS warning threshold : 0.0251 mW / -16.00 dBm
Modem lock state : OK
Lane 0
Tx power : 1.000 mW / 0.00 dBm
Module temperature : 0 degrees C / 32 degrees F
Rx power (total) : 0.000 mW / - Inf dBm
Rx power (signal) : 0.999 mW / -0.00 dBm
Lane chromatic dispersion : 6 ps/nm
Lane differential group delay : 3 ps
Lane Q2 factor : 15.40 dB
Lane carrier frequency offset : 0 MHz
Lane electrical SNR : 16.60 dB
Tx power high alarm : Off
Tx power low alarm : Off
Tx power high warning : Off
Tx power low warning : Off
Rx power high alarm : Off

```

```

Rx power low alarm : Off
Rx power high warning : Off
Rx power low warning : Off
Rx loss of signal alarm : Off
Wavelength unlocked alarm : Off
Laser end-of-life alarm : Off

Lane 1
Tx power : 1.000 mW / 0.00 dBm
Module temperature : 0 degrees C / 32 degrees F
Rx power (total) : 0.000 mW / - Inf dBm
Rx power (signal) : 0.999 mW / -0.00 dBm
Tx power high alarm : Off
Tx power low alarm : Off
Tx power high warning : Off
Tx power low warning : Off
Rx power high alarm : Off
Rx power low alarm : Off
Rx power high warning : Off
Rx power low warning : Off
Rx loss of signal alarm : Off
Wavelength unlocked alarm : Off
Laser end-of-life alarm : Off

Lane 2
Tx power : 1.000 mW / 0.00 dBm
Module temperature : 0 degrees C / 32 degrees F
Rx power (total) : 0.000 mW / - Inf dBm
Rx power (signal) : 0.999 mW / -0.00 dBm
Tx power high alarm : Off
Tx power low alarm : Off
Tx power high warning : Off
Tx power low warning : Off
Rx power high alarm : Off
Rx power low alarm : Off
Rx power high warning : Off
Rx power low warning : Off
Rx loss of signal alarm : Off
Wavelength unlocked alarm : Off
Laser end-of-life alarm : Off

Lane 3
Tx power : 1.000 mW / 0.00 dBm
Module temperature : 0 degrees C / 32 degrees F
Rx power (total) : 0.000 mW / - Inf dBm
Rx power (signal) : 0.999 mW / -0.00 dBm

```

```

Tx power high alarm : Off
Tx power low alarm : Off
Tx power high warning : Off
Tx power low warning : Off
Rx power high alarm : Off
Rx power low alarm : Off
Rx power high warning : Off
Rx power low warning : Off
Rx loss of signal alarm : Off
Wavelength unlocked alarm : Off
Laser end-of-life alarm : Off

```

### show interfaces diagnostics optics (for VCP)

```

user@host> show interfaces diagnostics optics vcp-2/0/1
Physical interface: vcp-2/0/1
 Laser bias current : 5.494 mA
 Laser output power : 0.2960 mW / -5.29 dBm
 Module temperature : 22 degrees C / 71 degrees F
 Module voltage : 3.2810 V
 Receiver signal average optical power : 0.2426 mW / -6.15 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : Off
 Laser rx power high warning : Off
 Laser rx power low warning : Off

```

```

Laser bias current high alarm threshold : 17.000 mA
Laser bias current low alarm threshold : 1.000 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6310 mW / -2.00 dBm
Laser output power low alarm threshold : 0.0670 mW / -11.74 dBm
Laser output power high warning threshold : 0.6310 mW / -2.00 dBm
Laser output power low warning threshold : 0.0790 mW / -11.02 dBm
Module temperature high alarm threshold : 95 degrees C / 203 degrees F
Module temperature low alarm threshold : -25 degrees C / -13 degrees F
Module temperature high warning threshold : 90 degrees C / 194 degrees F
Module temperature low warning threshold : -20 degrees C / -4 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7940 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

### show interfaces diagnostics optics (MPC7 with interfaces disabled)

```

user@host> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
 Module temperature : 34 degrees C / 93 degrees F
 Module voltage : 3.2660 V
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Module temperature high alarm threshold : 75 degrees C / 167 degrees F
 Module temperature low alarm threshold : -5 degrees C / 23 degrees F
 Module temperature high warning threshold : 70 degrees C / 158 degrees F
 Module temperature low warning threshold : 0 degrees C / 32 degrees F
 Module voltage high alarm threshold : 3.6300 V

```

```

Module voltage low alarm threshold : 2.9700 V
Module voltage high warning threshold : 3.4640 V
Module voltage low warning threshold : 3.1340 V
Laser bias current high alarm threshold : 9.999 mA
Laser bias current low alarm threshold : 0.499 mA
Laser bias current high warning threshold : 9.499 mA
Laser bias current low warning threshold : 0.999 mA
Laser output power high alarm threshold : 0.0000 mW / - Inf dBm
Laser output power low alarm threshold : 0.0000 mW / - Inf dBm
Laser output power high warning threshold : 0.0000 mW / - Inf dBm
Laser output power low warning threshold : 0.0000 mW / - Inf dBm
Laser rx power high alarm threshold : 2.1878 mW / 3.40 dBm
Laser rx power low alarm threshold : 0.0446 mW / -13.51 dBm
Laser rx power high warning threshold : 1.7378 mW / 2.40 dBm
Laser rx power low warning threshold : 0.1122 mW / -9.50 dBm

```

## Lane 0

```

Laser bias current : 6.697 mA
Laser output power : 0.738 mW / -1.32 dBm
Laser receiver power : 0.790 mW / -1.02 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off
Tx laser disabled alarm : Off

```

## Lane 1

```

Laser bias current : 6.961 mA
Laser output power : 0.908 mW / -0.42 dBm
Laser receiver power : 0.827 mW / -0.83 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off

```

```

Rx loss of signal alarm : Off
Tx laser disabled alarm : Off
Lane 2
Laser bias current : 6.926 mA
Laser output power : 0.888 mW / -0.51 dBm
Laser receiver power : 0.820 mW / -0.86 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off
Tx laser disabled alarm : Off
Lane 3
Laser bias current : 6.817 mA
Laser output power : 0.846 mW / -0.73 dBm
Laser receiver power : 0.827 mW / -0.82 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : Off
Laser receiver power high warning : Off
Laser receiver power low warning : Off
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : Off
Tx laser disabled alarm : Off

```

### show interfaces diagnostics optics for 100-Gigabit Ethernet (QSFP-100GE-DWDM2)

```

user@host> show interfaces diagnostics optics et-18/0/2
Physical interface: et-18/0/2
Module temperature : 37 degrees C / 98 degrees F
Module voltage : 3.2770 V
Grid Channel Number : 1 (191.40 THz)
Corrected Error Ratio : (6379 sec average) 3.29e-05

```

```

Uncorrected Words Ratio : (6379 sec average) 0.0e-05
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 : 79 degrees C / 174 degrees F
Module temperature low alarm threshold : -4 degrees C / 25 degrees F
Module temperature high warning threshold : 75 degrees C / 167 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.6300 V
Module voltage low alarm threshold : 2.9700 V
Module voltage high warning threshold : 3.4640 V
Module voltage low warning threshold : 3.1340 V
Laser bias current high alarm threshold : 109.999 mA
Laser bias current low alarm threshold : 19.999 mA
Laser bias current high warning threshold : 99.999 mA
Laser bias current low warning threshold : 29.999 mA
Laser output power high alarm threshold : 3.5481 mW / 5.50 dBm
Laser output power low alarm threshold : 0.2344 mW / -6.30 dBm
Laser output power high warning threshold : 2.8184 mW / 4.50 dBm
Laser output power low warning threshold : 0.2951 mW / -5.30 dBm
Laser rx power high alarm threshold : 3.5481 mW / 5.50 dBm
Laser rx power low alarm threshold : 0.0436 mW / -13.61 dBm
Laser rx power high warning threshold : 2.8183 mW / 4.50 dBm
Laser rx power low warning threshold : 0.0871 mW / -10.60 dBm
Lane 0
Laser bias current : 73.804 mA
Laser output power : 0.948 mW / -0.23 dBm
Laser receiver power : 0.003 mW / -25.23 dBm
Lane carrier frequency offset : 362 MHz
Lane SNR : 9.60 dB
TEC Current : 0.1 mA
Residual ISI : 125
PAM Histogram : 235
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off

```

```

Laser receiver power low alarm : On
Laser receiver power high warning : Off
Laser receiver power low warning : On
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : On
Tx laser disabled alarm : Off

Lane 1
Laser bias current : 83.363 mA
Laser output power : 1.042 mW / 0.18 dBm
Laser receiver power : 0.000 mW / - Inf dBm
Lane carrier frequency offset : 362 MHz
Lane SNR : 9.60 dB
TEC Current : 0.1 mA
Residual ISI : 125
PAM Histogram : 235
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser receiver power high alarm : Off
Laser receiver power low alarm : On
Laser receiver power high warning : Off
Laser receiver power low warning : On
Tx loss of signal functionality alarm : Off
Rx loss of signal alarm : On
Tx laser disabled alarm : Off

```

## Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 12.1 for PTX Series routers.

Command introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

## RELATED DOCUMENTATION

| *Determining Transceiver Support and Specifications*

## show interfaces (far-end-interval)

### IN THIS SECTION

- [Syntax | 1432](#)
- [Description | 1432](#)
- [Required Privilege Level | 1432](#)
- [Output Fields | 1432](#)
- [Sample Output | 1433](#)
- [Release Information | 1434](#)

### Syntax

```
show interfaces far-end-interval interface-fpc/pic/port
```

### Description

On channelized interfaces, display the far end interval data for the specified interface.

### Required Privilege Level

view

### Output Fields

[Table 152 on page 1432](#) lists the output fields for the **show interfaces far-end-interval** command. Output fields are listed in the approximate order in which they appear.

**Table 152: show interfaces far-end-interval Output Fields**

| Field Name         | Field Description              |
|--------------------|--------------------------------|
| Physical interface | Interface FPC/PIC/port values. |

**Table 152: show interfaces far-end-interval Output Fields (Continued)**

| Field Name          | Field Description                         |
|---------------------|-------------------------------------------|
| <b>SNMP ifIndex</b> | SNMP interface index value.               |
| <b>ES-L/P</b>       | Error detection—Errored seconds.          |
| <b>SES-L/P</b>      | Error detection—Severely errored seconds. |
| <b>UAS-L/P</b>      | Error detection—Unavailable seconds.      |

## Sample Output

### show interfaces far-end-interval coc12-5/2/0

```

user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
 05:30-current:
 ES-L: 1, SES-L: 1, UAS-L: 0
 05:15-05:30:
 ES-L: 0, SES-L: 0, UAS-L: 0
 05:00-05:15:
 ES-L: 0, SES-L: 0, UAS-L: 0
 04:45-05:00:
 ES-L: 0, SES-L: 0, UAS-L: 0
 04:30-04:45:
 ES-L: 0, SES-L: 0, UAS-L: 0
 04:15-04:30:
 ES-L: 0, SES-L: 0, UAS-L: 0
 04:00-04:15:
 ...

```

## show interfaces far-end-interval coc1-5/2/1:1

```

user@host> run show interfaces far-end-interval coc1-5/2/1:1
Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
05:30-current:
 ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:15-05:30:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:00-05:15:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:45-05:00:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:30-04:45:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:15-04:30:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:00-04:15:

```

## Release Information

Command introduced in Junos OS Release 9.4.

## show interfaces (Fast Ethernet)

### IN THIS SECTION

- [Syntax | 1435](#)
- [Description | 1435](#)
- [Options | 1435](#)
- [Required Privilege Level | 1435](#)
- [Output Fields | 1435](#)
- [Sample Output | 1460](#)
- [Release Information | 1463](#)

## Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Description

Display status information about the specified Fast Ethernet interface.

## Options

|                                           |                                                                                         |
|-------------------------------------------|-----------------------------------------------------------------------------------------|
| <i>interface-type</i>                     | On M Series and T Series routers, the interface type is <b><i>fe-fpc/pic/port</i></b> . |
| <b>brief   detail   extensive   terse</b> | (Optional) Display the specified level of output.                                       |
| <b>descriptions</b>                       | (Optional) Display interface description strings.                                       |
| <b>media</b>                              | (Optional) Display media-specific information about network interfaces.                 |
| <b>snmp-index <i>snmp-index</i></b>       | (Optional) Display information for the specified SNMP index of the interface.           |
| <b>statistics</b>                         | (Optional) Display static interface statistics.                                         |

## Required Privilege Level

view

## Output Fields

[Table 153 on page 1436](#) lists the output fields for the **show interfaces** (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 153: show interfaces Fast Ethernet Output Fields

| Field Name                | Field Description                                                                                                                             | Level of Output                 |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Physical Interface</b> |                                                                                                                                               |                                 |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                               | All levels                      |
| <b>Enabled</b>            | State of the interface. Possible values are described in the “Enabled Field” section under <a href="#">Common Output Fields Description</a> . | All levels                      |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                           | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                 | <b>detail extensive</b><br>none |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                             | <b>detail extensive</b>         |
| <b>Link-level type</b>    | Encapsulation being used on the physical interface.                                                                                           | All levels                      |
| <b>MTU</b>                | Maximum transmission unit size on the physical interface.                                                                                     | All levels                      |
| <b>Link-mode</b>          | Type of link connection configured for the physical interface:<br><b>Full-duplex</b> or <b>Half-duplex</b>                                    | <b>extensive</b>                |
| <b>Speed</b>              | Speed at which the interface is running.                                                                                                      | All levels                      |
| <b>Loopback</b>           | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                | All levels                      |
| <b>Source filtering</b>   | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                  | All levels                      |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                   | Level of Output |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>LAN-PHY mode</b>     | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.                                                       | All levels      |
| <b>WAN-PHY mode</b>     | 10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.                           | All levels      |
| <b>Unidirectional</b>   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.                                                                   | All levels      |
| <b>Flow control</b>     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                            | All levels      |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                           | All levels      |
| <b>Remote-fault</b>     | (Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul> | All levels      |
| <b>Device flags</b>     | Information about the physical device. Possible values are described in the “Device Flags” section under <a href="#">Common Output Fields Description</a> .                                                                                         | All levels      |
| <b>Interface flags</b>  | Information about the interface. Possible values are described in the “Interface Flags” section under <a href="#">Common Output Fields Description</a> .                                                                                            | All levels      |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                           | Level of Output                 |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Link flags</b>       | Information about the link. Possible values are described in the "Links Flags" section under <a href="#">Common Output Fields Description</a> .                                                                                                             | All levels                      |
| <b>Wavelength</b>       | (10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).                                                                                                                      | All levels                      |
| <b>Frequency</b>        | (10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                            | All levels                      |
| <b>CoS queues</b>       | Number of CoS queues configured.                                                                                                                                                                                                                            | <b>detail extensive</b><br>none |
| <b>Schedulers</b>       | (GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.                                                                                                                                                          | <b>extensive</b>                |
| <b>Hold-times</b>       | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                         | <b>detail extensive</b>         |
| <b>Current address</b>  | Configured MAC address.                                                                                                                                                                                                                                     | <b>detail extensive</b><br>none |
| <b>Hardware address</b> | Hardware MAC address.                                                                                                                                                                                                                                       | <b>detail extensive</b><br>none |
| <b>Last flapped</b>     | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive</b><br>none |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | None specified          |
| <b>Output Rate</b>             | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | None specified          |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the <a href="#">show interfaces</a> command.</p> | <b>detail extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b> | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame terminations and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output  |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Output errors</b> | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminations and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output         |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
|                                | <ul style="list-style-type: none"> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                         |
| <b>Egress queues</b>           | <p>Total number of egress queues supported on the specified interface.</p> <p><b>NOTE:</b> In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the <b>show interfaces</b> command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p> | <b>detail extensive</b> |
| <b>Queue counters (Egress)</b> | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                        |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| <b>Ingress queues</b>                   | Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b>                       |
| <b>Queue counters (Ingress)</b>         | CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>                       |
| <b>Active alarms and Active defects</b> | Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b> . <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul> | <b>detail extensive</b><br><b>none</b> |
| <b>OTN FEC statistics</b>               | The forward error correction (FEC) counters provide the following statistics: <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—The count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul>                                                                                                                                                                                                                                                                                                                                                 |                                        |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output         |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>PCS statistics</b> | <p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> <li>• <b>Bit errors</b>—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.</li> <li>• <b>Errored blocks</b>—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.</li> </ul> | <b>detail extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output  |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>MAC statistics</b> | <p><b>Receive</b> and <b>Transmit</b> statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the <a href="#">show interfaces</a> command.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of frames that exceed 1518 octets.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output  |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                       | <ul style="list-style-type: none"> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”</li> </ul> |                  |
| <b>OTN Received Overhead Bytes</b>    | APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b> |
| <b>OTN Transmitted Overhead Bytes</b> | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output  |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Filter statistics</b> | <p><b>Receive</b> and <b>Transmit</b> statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> <li>• <b>Input packet count</b>—Number of packets received from the MAC hardware that the filter processed.</li> <li>• <b>Input packet rejects</b>—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address.</li> <li>• <b>Input DA rejects</b>—Number of packets that the filter rejected because the destination 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 routing device from the far-end system, either there is a bad 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 routing device (which the routing device is rejecting).</li> <li>• <b>Input SA rejects</b>—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect.</li> <li>• <b>Output packet count</b>—Number of packets that the filter has given to the MAC hardware.</li> <li>• <b>Output packet pad count</b>—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done</li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output  |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                | <p>only on small ARP packets, but some very small 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.</p> <ul style="list-style-type: none"> <li>• <b>Output packet error count</b>—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.</li> <li>• <b>CAM destination filters, CAM source filters</b>—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.</li> </ul> |                  |
| <b>PMA PHY</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>PHY Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>                                                                                                                                                                                                                                                                                                                                   | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (*Continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>WIS section</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B1</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOL</b>—Loss of light</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (*Continued*)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output  |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>WIS line</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. State other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output  |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>WIS path</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output  |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Autonegotiation information</b> | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—<b>OK</b> when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> </ul> | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output  |
|--------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                                        | <ul style="list-style-type: none"> <li>• <b>Local resolution</b>—Information from the link partner:               <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> |                  |
| <b>Received path trace,<br/>Transmitted path trace</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other routing device manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>                                                                                                                                                                                                           | <b>extensive</b> |
| <b>Packet Forwarding Engine configuration</b>          | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output                 |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>CoS information</b>   | <p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>                |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                 |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels                      |
| <b>Index</b>             | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>      | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b><br>none |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name        | Field Description                                                                                                                                                        | Level of Output         |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Generation</b> | Unique number for use by Juniper Networks technical support only.                                                                                                        | <b>detail extensive</b> |
| <b>Flags</b>      | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <a href="#">Common Output Fields Description</a> . | All levels              |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                                       |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|
| VLAN-Tag   | <p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li>• <b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information.</li> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li>• <b>swap-push</b>—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</li> <li>• <b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value.</li> <li>• <b>pop-swap</b>—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</li> <li>• <b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul> | <p><b>brief detail</b><br/> <b>extensive none</b></p> |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                        | Level of Output                 |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Demux:</b>                  | <p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> <li>• Source Family Inet</li> <li>• Destination Family Inet</li> </ul>                                                                                                           | <b>detail extensive</b><br>none |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                  | All levels                      |
| <b>Protocol</b>                | Protocol family. Possible values are described in the “Protocol Field” section under <a href="#">Common Output Fields Description</a> .                                                                                                                                                                                                                                  | <b>detail extensive</b><br>none |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b><br>none |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                          | <b>detail extensive</b><br>none |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul> | <b>detail extensive</b>         |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                       | <b>extensive</b>                |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to the routing device.                                                                                                                                                                                                                                                                                                     | <b>extensive</b>                |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Transit statistics</b>       | <p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p> | <b>extensive</b>                |
| <b>Generation</b>               | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |
| <b>Route Table</b>              | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b><br>none |
| <b>Flags</b>                    | Information about protocol family flags. Possible values are described in the “Family Flags” section under <a href="#">Common Output Fields Description</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b>         |
| <b>Donor interface</b>          | (Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b><br>none |
| <b>Preferred source address</b> | (Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b><br>none |
| <b>Input Filters</b>            | Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>         |

Table 153: show interfaces Fast Ethernet Output Fields (Continued)

| Field Name                   | Field Description                                                                                                                                                                                    | Level of Output                 |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Output Filters</b>        | Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces. | <b>detail extensive</b>         |
| <b>Mac-Validate Failures</b> | Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.                                           | <b>detail extensive</b><br>none |
| <b>Addresses, Flags</b>      | Information about the address flags. Possible values are described in the “Addresses Flags” section under <a href="#">Common Output Fields Description</a> .                                         | <b>detail extensive</b><br>none |
| <i>protocol-family</i>       | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                             | <b>brief</b>                    |
| <b>Flags</b>                 | Information about address flag (possible values are described in the “Addresses Flags” section under <a href="#">Common Output Fields Description</a> ).                                             | <b>detail extensive</b><br>none |
| <b>Destination</b>           | IP address of the remote side of the connection.                                                                                                                                                     | <b>detail extensive</b><br>none |
| <b>Local</b>                 | IP address of the logical interface.                                                                                                                                                                 | <b>detail extensive</b><br>none |
| <b>Broadcast</b>             | Broadcast address of the logical interlace.                                                                                                                                                          | <b>detail extensive</b><br>none |
| <b>Generation</b>            | Unique number for use by Juniper Networks technical support only.                                                                                                                                    | <b>detail extensive</b>         |

## Sample Output

### show interfaces (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 22
 Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 4 supported, 4 maximum usable queues
 Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
 Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None
 Active defects : None
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500
 Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
```

### show interfaces brief (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
Logical interface fe-0/0/0.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet 203.0.113.1/24
```

**show interfaces detail (Fast Ethernet)**

```

user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 22, Generation: 5391
 Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 4 supported, 4 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:3f:38
 Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 42 0 bps
 Input packets : 0 0 pps
 Output packets: 1 0 pps
 Active alarms : None
 Active defects : None
 Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500, Generation: 105, Route table: 0
 Flags: Is-Primary, Mac-Validate-Strict
 Mac-Validate Failures: Packets: 0, Bytes: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
 Generation: 136

```

**show interfaces extensive (Fast Ethernet)**

```

user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 22, Generation: 5391
 Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
 100mbps, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000

```

```

CoS queues : 4 supported, 4 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 42 0 bps
 Input packets: 0 0 pps
 Output packets: 1 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: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms : None
Active defects : None
MAC statistics:

```

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 64       |
| Total packets      | 0       | 1        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 1        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |
| VLAN tagged frames | 0       |          |
| Code violations    | 0       |          |

```

Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 1
 Output packet pad count 0
 Output packet error count 0
CAM destination filters: 1, CAM source filters: 0

```

```

Autonegotiation information:
 Negotiation status: Complete
 Link partner:
 Link partner: Full-duplex, Flow control: None, Remote fault: Ok
 Local resolution:
Packet Forwarding Engine configuration:
 Destination slot: 0
CoS information:
 Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 500000000 5 0 low none
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
 Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
 Generation: 136

```

## Release Information

Command introduced before Junos OS Release 7.4.

## show interfaces

### IN THIS SECTION

- [Syntax \(Gigabit Ethernet\) | 1464](#)
- [Syntax \(10 Gigabit Ethernet\) | 1464](#)
- [Syntax \(ACX5448, ACX5448-D, ACX710\) | 1464](#)
- [Syntax \(QFX5130-32CD\) | 1465](#)
- [Syntax \(SRX Series Devices and \(vSRX and vSRX 3.0 platforms\)\) | 1465](#)
- [Description | 1466](#)
- [Options | 1466](#)
- [Additional Information | 1468](#)

- [Required Privilege Level | 1468](#)
- [Output Fields | 1468](#)
- [Sample output for G.fast and Annex J support | 1538](#)
- [Sample Output Gigabit Ethernet | 1554](#)
- [Sample Output | 1567](#)
- [Release Information | 1601](#)

## Syntax (Gigabit Ethernet)

```
show interfaces ge-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Syntax (10 Gigabit Ethernet)

```
show interfaces xe-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Syntax (ACX5448, ACX5448-D, ACX710)

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Syntax (QFX5130-32CD)

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Syntax (SRX Series Devices and (vSRX and vSRX 3.0 platforms))

```
show interfaces (
 <interface-name>
 <brief | detail | extensive | terse>
 <controller interface-name>|
 <descriptions interface-name>|
 <destination-class (all | destination-class-name logical-interface-name)>|
 <diagnostics optics interface-name>|
 <far-end-interval interface-fpc/pic/port>|
 <filters interface-name>|
 <flow-statistics interface-name>|
 <interval interface-name>|
 <load-balancing (detail | interface-name)>|
 <mac-database mac-address mac-address>|
 <mc-ae id identifier unit number revertive-info>|
 <media interface-name>|
 <policers interface-name>|
 <queue both-ingress-egress egress forwarding-class forwarding-class ingress
 12-statistics>|
 <redundancy (detail | interface-name)>|
 <routing brief detail summary interface-name>|
 <routing-instance (all | instance-name)>|
 <snmp-index snmp-index>|
 <source-class (all | destination-class-name logical-interface-name)>|
 <statistics interface-name>|
 <switch-port switch-port number>|
 <transport pm (all | optics | otn) (all | current | currentday | interval |
 previousday) (all | interface-name)>|
 <zone interface-name>|
```

```
<dsl-sfp-options (adsl-options | gfast-options | vdsl-options) >
)
```

## Description

Display status information about the specified Gigabit Ethernet interface.

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series, vSRX, and vSRX 3.0 platforms running Junos OS.

SRX4600 supports 40-Gigabit Ethernet breakouts only in PIC mode. Use the **show interfaces extensive** command to view the speed configured for the interface on SRX4600. Reboot the device for the changed configuration to take effect.

On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Starting in Junos OS Release 18.4R1, Output fields **Next-hop** and **vpls-status** is displayed in the **show interfaces *interface name* detail** command, only for Layer 2 protocols on MX480 routers.

In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the **show interfaces *interface-name*<media><extensive>** command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed. In the earlier Junos OS releases, incorrect autonegotiation status was displayed even when autonegotiation was disabled.

QFX5130-32CD switches does not display the **Filters statistics** when the **show interfaces extensive** command is executed due to interface-level filter statistics related hardware limitations. See "[show interfaces extensive \(QFX5130-32CD\)](#)" on page 1598.

Starting in Junos OS Release 20.4R1, we support G.fast and Annex J specification with SFP xDSL for ADSL2/ADSL2+ and all VDSL2 profiles on SRX380, SRX300, SRX320, SRX340, and SRX345 devices.

## Options

For Gigabit interfaces:

**ge-*fpc*/*pic*/*port***            Display standard information about the specified Gigabit Ethernet interface.

**NOTE:** Interfaces with different speeds are named uniformly with ge-0/0/x for backward compatibility. Use the **show interfaces** command to view the interface speeds.

|                                           |                                                                               |
|-------------------------------------------|-------------------------------------------------------------------------------|
| <b>brief   detail   extensive   terse</b> | (Optional) Display the specified level of output.                             |
| <b>descriptions</b>                       | (Optional) Display interface description strings.                             |
| <b>media</b>                              | (Optional) Display media-specific information about network interfaces.       |
| <b>snmp-index <i>snmp-index</i></b>       | (Optional) Display information for the specified SNMP index of the interface. |
| <b>statistics</b>                         | (Optional) Display static interface statistics.                               |

For 10 Gigabit interfaces:

|                                           |                                                                                 |
|-------------------------------------------|---------------------------------------------------------------------------------|
| <b><i>xe-fpc/pic/port</i></b>             | Display standard information about the specified 10-Gigabit Ethernet interface. |
| <b>brief   detail   extensive   terse</b> | (Optional) Display the specified level of output.                               |
| <b>descriptions</b>                       | (Optional) Display interface description strings.                               |
| <b>media</b>                              | (Optional) Display media-specific information about network interfaces.         |
| <b>snmp-index <i>snmp-index</i></b>       | (Optional) Display information for the specified SNMP index of the interface.   |
| <b>statistics</b>                         | (Optional) Display static interface statistics.                                 |

For SRX interfaces:

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>interface-name</b> | (Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number. <ul style="list-style-type: none"> <li>• <b><i>at-pim/0/port</i></b>—ATM-over-ADSL or ATM-over-SHDSL interface.</li> <li>• <b><i>ce1-pim/0/port</i></b>—Channelized E1 interface.</li> <li>• <b><i>cl-0/0/8</i></b>—3G wireless modem interface for SRX320 devices.</li> <li>• <b><i>ct1-pim/0/port</i></b>—Channelized T1 interface.</li> </ul> |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

- **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
- **e1-*pim/0/port***—E1 interface.
- **e3-*pim/0/port***—E3 interface.
- **fe-*pim/0/port***—Fast Ethernet interface.
- **ge-*pim/0/port***—Gigabit Ethernet interface.
- **se-*pim/0/port***—Serial interface.
- **t1-*pim/0/port***—T1 (also called DS1) interface.
- **t3-*pim/0/port***—T3 (also called DS3) interface.
- **wx-*slot/0/0***—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

## Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

## Required Privilege Level

view

## Output Fields

[Table 154 on page 1468](#) describes the output fields for the **show interfaces** (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see [Table 155 on page 1523](#).

**Table 154: show interfaces (Gigabit Ethernet) Output Fields**

| Field Name                | Field Description               | Level of Output |
|---------------------------|---------------------------------|-----------------|
| <b>Physical Interface</b> |                                 |                 |
| <b>Physical interface</b> | Name of the physical interface. | All levels      |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                             | Level of Output                 |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Enabled</b>          | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .                                                          | All levels                      |
| <b>Interface index</b>  | Index number of the physical interface, which reflects its initialization sequence.                                                                                                           | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>     | SNMP index number for the physical interface.                                                                                                                                                 | <b>detail extensive</b><br>none |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                                                             | <b>detail extensive</b>         |
| <b>Link-level type</b>  | Encapsulation being used on the physical interface.                                                                                                                                           | All levels                      |
| <b>MTU</b>              | Maximum transmission unit size on the physical interface.                                                                                                                                     | All levels                      |
| <b>Speed</b>            | Speed at which the interface is running.                                                                                                                                                      | All levels                      |
| <b>Loopback</b>         | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                | All levels                      |
| <b>Source filtering</b> | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                  | All levels                      |
| <b>LAN-PHY mode</b>     | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications. | All levels                      |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                   | Level of Output |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>WAN-PHY mode</b>     | 10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.                           | All levels      |
| <b>Unidirectional</b>   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.                                                                   | All levels      |
| <b>Flow control</b>     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                            | All levels      |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                           | All levels      |
| <b>Remote-fault</b>     | (Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul> | All levels      |
| <b>Device flags</b>     | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .                                                                                                  | All levels      |
| <b>Interface flags</b>  | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                     | All levels      |
| <b>Link flags</b>       | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .                                                                                                              | All levels      |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (*Continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                           | Level of Output                 |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Wavelength</b>       | (10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).                                                                                                                      | All levels                      |
| <b>Frequency</b>        | (10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                            | All levels                      |
| <b>CoS queues</b>       | Number of CoS queues configured.                                                                                                                                                                                                                            | <b>detail extensive</b><br>none |
| <b>Schedulers</b>       | (Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.                                                                                                                                                         | <b>extensive</b>                |
| <b>Hold-times</b>       | Current interface hold-time up and hold-time down, in milliseconds (ms).                                                                                                                                                                                    | <b>detail extensive</b>         |
| <b>Current address</b>  | Configured MAC address.                                                                                                                                                                                                                                     | <b>detail extensive</b><br>none |
| <b>Hardware address</b> | Hardware MAC address.                                                                                                                                                                                                                                       | <b>detail extensive</b><br>none |
| <b>Last flapped</b>     | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive</b><br>none |

Table 154: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

| Field Name                      | Field Description                                                                                                                                                                                                                                                           | Level of Output         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Input Rate</b>               | Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. | None                    |
| <b>Output Rate</b>              | Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.                                        | None                    |
| <b>Statistics last cleared</b>  | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>Egress account overhead</b>  | Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.                                                                                                                                                                                 | <b>detail extensive</b> |
| <b>Ingress account overhead</b> | Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.                                                                                                                                                                                | <b>detail extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output         |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the <a href="#">"show interfaces " on page 1463</a> command.</p> | <b>detail extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b> | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame terminated and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output  |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Output errors</b> | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminated and FCS errors.</li> <li>• <b>Drops</b>—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.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Drops</b> field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number must always be 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output         |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
|                      | <ul style="list-style-type: none"> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                         |
| <b>Egress queues</b> | <p>Total number of egress queues supported on the specified interface.</p> <p><b>NOTE:</b> In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the <b>show interfaces</b> command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p> | <b>detail extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (*Continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Queue counters (Egress)</b>  | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Dropped packets</b> field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> | <b>detail extensive</b> |
| <b>Ingress queues</b>           | Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b>        |
| <b>Queue counters (Ingress)</b> | <p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                        | <b>extensive</b>        |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output                 |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul> | <b>detail extensive</b><br>none |
| Interface transmit statistics           | <p>(On MX Series devices) Status of the <b>interface-transmit-statistics</b> configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—When the <b>interface-transmit-statistics</b> statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface.</li> <li>• <b>Disabled</b>—When the <b>interface-transmit-statistics</b> statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface.</li> </ul>                                                                                     | <b>detail extensive</b>         |
| <b>OTN FEC statistics</b>               | <p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—Count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul>                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>         |

Table 154: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output         |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>PCS statistics</b> | <p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> <li>• <b>Bit errors</b>—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.</li> <li>• <b>Errored blocks</b>—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.</li> </ul> | <b>detail extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output         |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Link Degrade</b> | <p>Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.</p> <ul style="list-style-type: none"> <li>• <b>Link Monitoring</b>—Indicates if physical link degrade monitoring is enabled on the interface. <ul style="list-style-type: none"> <li>• <b>Enable</b>—Indicates that link degrade monitoring has been enabled (using the <b>link-degrade-monitor</b> statement) on the interface.</li> <li>• <b>Disable</b>—Indicates that link degrade monitoring has not been enabled on the interface. If link degrade monitoring has not been enabled, the output does not show any related information, such as BER values and thresholds.</li> </ul> </li> <li>• <b>Link Degrade Set Threshold</b>—The BER threshold value at which the link is considered degraded and a corrective action is triggered.</li> <li>• <b>Link Degrade Clear Threshold</b>—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted.</li> <li>• <b>Estimated BER</b>—The estimated bit error rate.</li> <li>• <b>Link-degrade event</b>—Shows link degrade event information. <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Time (in seconds) elapsed after a link degrade event occurred.</li> <li>• <b>Count</b>—The number of link degrade events recorded.</li> <li>• <b>State</b>—Shows the link degrade status (example: <b>Defect Active</b>).</li> </ul> </li> </ul> | <b>detail extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output  |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>MAC statistics</b> | <p><b>Receive</b> and <b>Transmit</b> statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the <a href="#">"show interfaces "</a> on page 1463 command.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> <li>• Packet length exceeds interface MTU, or</li> <li>• Packet length exceeds MRU</li> </ul> </li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4</li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output  |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                       | <p>(10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</p> <ul style="list-style-type: none"> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.</li> </ul> <p><b>NOTE:</b> The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the <b>VLAN tagged frames</b> field displays 0 when the <b>show interfaces</b> command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> <ul style="list-style-type: none"> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”</li> </ul> |                  |
| <b>OTN Received Overhead Bytes</b>    | APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b> |
| <b>OTN Transmitted Overhead Bytes</b> | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output  |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Filter statistics</b> | <p><b>Receive</b> and <b>Transmit</b> statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected.</p> <ul style="list-style-type: none"> <li>• <b>Input packet count</b>—Number of packets received from the MAC hardware that the filter processed.</li> <li>• <b>Input packet rejects</b>—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address.</li> <li>• <b>Input DA rejects</b>—Number of packets that the filter rejected because the destination 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 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).</li> <li>• <b>Input SA rejects</b>—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field must increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect.</li> <li>• <b>Output packet count</b>—Number of packets that the filter has given to the MAC hardware.</li> <li>• <b>Output packet pad count</b>—Number of packets 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 IP packets</li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output  |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                | <p>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.</p> <ul style="list-style-type: none"> <li>• <b>Output packet error count</b>—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field must not increment.</li> <li>• <b>CAM destination filters, CAM source filters</b>—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0.</li> </ul> |                  |
| <b>PMA PHY</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>PHY Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>                                                                                                                                                                                                                                                                     | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (*Continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>WIS section</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B1</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOL</b>—Loss of light</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output  |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>WIS line</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output  |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>WIS path</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul> <p>Subfields are:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output  |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Autonegotiation information</b> | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—<b>OK</b> when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner</b>—Information from the remote Ethernet device: <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the link partner, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), <b>Symmetric/Asymmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> on transmit), and <b>None</b> (link partner does not support flow control).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> </ul> | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output  |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                                    | <ul style="list-style-type: none"> <li>• <b>Local resolution</b>—Information from the local Ethernet device:               <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are <b>Symmetric/Asymmetric</b> (local device supports <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> on receive) and <b>None</b> (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display <b>Symmetric</b> (local device supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (local device supports <b>PAUSE</b> on receive), and <b>None</b> (local device does not support flow control).</li> </ul> </li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> |                  |
| <b>Received path trace, Transmitted path trace</b> | (10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>extensive</b> |
| <b>Packet Forwarding Engine configuration</b>      | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>extensive</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output                 |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>CoS information</b>   | <p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>                |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                 |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | All levels                      |
| <b>Index</b>             | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>      | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b><br>none |
| <b>Generation</b>        | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>         |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output                              |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <b>Flags</b>    | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                                   |
| <b>VLAN-Tag</b> | <p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li>• <b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information.</li> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li>• <b>swap-push</b>—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</li> <li>• <b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value.</li> <li>• <b>pop-swap</b>—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</li> <li>• <b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul> | <b>brief detail</b><br><b>extensive none</b> |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output                              |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <b>Demux</b>         | <p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> <li>• Source Family Inet</li> <li>• Destination Family Inet</li> </ul>                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b><br>none              |
| <b>Encapsulation</b> | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | All levels                                   |
| <b>ACI VLAN</b>      | <p>Information displayed for agent circuit identifier (ACI) interface set configured with the <b>agent-circuit-id</b> autoconfiguration stanza.</p> <p><b>Dynamic Profile</b>—Name of the dynamic profile that defines the ACI interface set.</p> <p>If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.</p> <p><b>NOTE:</b> The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the <b>line-identity</b> autoconfiguration stanza.</p> | <b>brief detail</b><br><b>extensive</b> none |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output                 |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| Line Identity   | <p>Information displayed for access-line-identifier (ALI) interface sets configured with the <b>line-identity</b> autoconfiguration stanza.</p> <ul style="list-style-type: none"> <li>• <b>Dynamic Profile</b>—Name of the dynamic profile that defines the ALI interface set.</li> <li>• Trusted option used to create the ALI interface set: <b>Circuit-id</b>, <b>Remote-id</b>, or <b>Accept-no-ids</b>. More than one option can be configured.</li> </ul> <p>If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information.</p> <p><b>NOTE:</b> The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the <b>agent-circuit-id</b> autoconfiguration stanza.</p> | <b>detail</b>                   |
| <b>Protocol</b> | Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b><br>none |
| <b>MTU</b>      | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b><br>none |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                                               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                 |
|----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Neighbor Discovery Protocol (NDP)Queue Statistics</b> | <p>NDP statistics for protocol <b>inet6</b> under logical interface statistics.</p> <ul style="list-style-type: none"> <li>• <b>Max nh cache</b>—Maximum interface neighbor discovery nexthop cache size.</li> <li>• <b>New hold nh limit</b>—Maximum number of new unresolved nexthops.</li> <li>• <b>Curr nh cnt</b>—Current number of resolved nexthops in the NDP queue.</li> <li>• <b>Curr new hold cnt</b>—Current number of unresolved nexthops in the NDP queue.</li> <li>• <b>NH drop cnt</b>—Number of NDP requests not serviced.</li> </ul> | All levels                      |
| <b>Dynamic Profile</b>                                   | Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b><br>none |
| <b>Service Name Table</b>                                | Name of the service name table for the interface configured with a PPPoE family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b><br>none |
| <b>Max Sessions</b>                                      | Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b><br>none |
| <b>Duplicate Protection</b>                              | State of PPPoE duplicate protection: <b>On</b> or <b>Off</b> . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.                                                                                                                                                                                                                                                                             | <b>detail extensive</b><br>none |
| <b>Direct Connect</b>                                    | State of the configuration to ignore DSL Forum VSAs: <b>On</b> or <b>Off</b> . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b><br>none |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output                 |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>AC Name</b>                 | Name of the access concentrator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b><br>none |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b><br>none |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul> | <b>detail extensive</b>         |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>                |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>extensive</b>                |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (*Continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Transit statistics</b>       | <p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p> | <b>extensive</b>                |
| <b>Generation</b>               | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |
| <b>Route Table</b>              | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b><br>none |
| <b>Flags</b>                    | Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>         |
| <b>Donor interface</b>          | (Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b><br>none |
| <b>Preferred source address</b> | (Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive</b><br>none |
| <b>Input Filters</b>            | Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail extensive</b>         |

Table 154: show interfaces (Gigabit Ethernet) Output Fields (Continued)

| Field Name                   | Field Description                                                                                                                                                                                    | Level of Output                 |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Output Filters</b>        | Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces. | <b>detail extensive</b>         |
| <b>Mac-Validate Failures</b> | Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.                                           | <b>detail extensive</b><br>none |
| <b>Addresses, Flags</b>      | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                                                  | <b>detail extensive</b><br>none |
| <i>protocol-family</i>       | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                             | <b>brief</b>                    |
| <b>Flags</b>                 | Information about the address flag. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                                                   | <b>detail extensive</b><br>none |
| <b>Destination</b>           | IP address of the remote side of the connection.                                                                                                                                                     | <b>detail extensive</b><br>none |
| <b>Local</b>                 | IP address of the logical interface.                                                                                                                                                                 | <b>detail extensive</b><br>none |
| <b>Broadcast</b>             | Broadcast address of the logical interface.                                                                                                                                                          | <b>detail extensive</b><br>none |
| <b>Generation</b>            | Unique number for use by Juniper Networks technical support only.                                                                                                                                    | <b>detail extensive</b>         |

The following table describes the output fields for the **show interfaces** (10-Gigabit Ethernet) command.

| Field Name                | Field Description                                                                                                                                                                             | Level of Output                           |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <b>Physical interface</b> | Name of the physical interface.                                                                                                                                                               | All levels                                |
| <b>Enabled</b>            | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .                                                          | All levels                                |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                                                                           | <b>detail</b><br><b>extensive</b><br>none |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                                                                 | <b>detail</b><br><b>extensive</b><br>none |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                                                                             | <b>detail</b><br><b>extensive</b>         |
| <b>Link-level type</b>    | Encapsulation being used on the physical interface.                                                                                                                                           | All levels                                |
| <b>MTU</b>                | Maximum transmission unit size on the physical interface.                                                                                                                                     | All levels                                |
| <b>Speed</b>              | Speed at which the interface is running.                                                                                                                                                      | All levels                                |
| <b>Loopback</b>           | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .                                                                | All levels                                |
| <b>Source filtering</b>   | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                  | All levels                                |
| <b>LAN-PHY mode</b>       | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications. | All levels                                |

|                         |                                                                                                                                                                                                                                                     |            |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| <b>WAN-PHY mode</b>     | 10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.                           | All levels |
| <b>Unidirectional</b>   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.                                                                   | All levels |
| <b>Flow control</b>     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                                                            | All levels |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                           | All levels |
| <b>Remote-fault</b>     | (Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul> | All levels |
| <b>Device flags</b>     | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .                                                                                                  | All levels |
| <b>Interface flags</b>  | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                     | All levels |
| <b>Link flags</b>       | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .                                                                                                              | All levels |
| <b>Wavelength</b>       | (10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).                                                                                                              | All levels |

|                         |                                                                                                                                                                                                                                                                             |                             |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| <b>Frequency</b>        | (10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).                                                                                                                                            | All levels                  |
| <b>CoS queues</b>       | Number of CoS queues configured.                                                                                                                                                                                                                                            | detail<br>extensive<br>none |
| <b>Schedulers</b>       | (Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only)<br>Number of CoS schedulers configured.                                                                                                                                                                      | extensive                   |
| <b>Hold-times</b>       | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                         | detail<br>extensive         |
| <b>Current address</b>  | Configured MAC address.                                                                                                                                                                                                                                                     | detail<br>extensive<br>none |
| <b>Hardware address</b> | Hardware MAC address.                                                                                                                                                                                                                                                       | detail<br>extensive<br>none |
| <b>Last flapped</b>     | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                 | detail<br>extensive<br>none |
| <b>Input Rate</b>       | Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. | None<br>specified           |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| <b>Output Rate</b>              | Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | None specified              |
| <b>Statistics last cleared</b>  | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail<br/>extensive</b> |
| <b>Egress account overhead</b>  | Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail<br/>extensive</b> |
| <b>Ingress account overhead</b> | Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail<br/>extensive</b> |
| <b>Traffic statistics</b>       | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail<br/>extensive</b> |

**Input errors**

Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:

- **Errors**—Sum of the incoming frame terminated and FCS errors.
- **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.
- **Framing errors**—Number of packets received with an invalid frame checksum (FCS).
- **Runts**—Number of frames received that are smaller than the runt threshold.
- **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.
- **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement.
- **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame.
- **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.
- **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.
- **Resource errors**—Sum of transmit drops.

**extensive**

**Output errors**

Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:

- **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.
- **Errors**—Sum of the outgoing frame terminated and FCS errors.
- **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.
- **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.
- **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.
- **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.
- **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
- **MTU errors**—Number of packets whose size exceeded the MTU of the interface.
- **Resource errors**—Sum of transmit drops.

**extensive**

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                             |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| <b>Egress queues</b>           | <p>Total number of egress queues supported on the specified interface.</p> <p><b>NOTE:</b> In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the <b>show interfaces</b> command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p> | <b>detail<br/>extensive</b> |
| <b>Queue counters (Egress)</b> | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail<br/>extensive</b> |
| <b>Ingress queues</b>          | <p>Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>extensive</b>            |

|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                           |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <b>Queue counters (Ingress)</b>         | <p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                              | <b>extensive</b>                          |
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul> | <b>detail</b><br><b>extensive</b><br>none |
| <b>OTN alarms</b>                       | Active OTN alarms identified on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b><br><b>extensive</b>         |
| <b>OTN defects</b>                      | OTN defects received on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail</b><br><b>extensive</b>         |
| <b>OTN FEC Mode</b>                     | <p>The FECmode configured on the interface.</p> <ul style="list-style-type: none"> <li>• <b>efec</b>—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors.</li> <li>• <b>gfec</b>—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors.</li> <li>• <b>none</b>—FEC mode is not configured.</li> </ul>                                                                                                                                                                                                                                                                                                                                 | <b>detail</b><br><b>extensive</b>         |

|                    |                                                                                                                                                                                                                                                                                                                                                    |                     |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| OTN Rate           | <p>OTN mode.</p> <ul style="list-style-type: none"> <li>• <b>fixed-stuff-bytes</b>—Fixed stuff bytes 11.0957 Gbps.</li> <li>• <b>no-fixed-stuff-bytes</b>—No fixed stuff bytes 11.0491 Gbps.</li> <li>• <b>pass-through</b>—Enable OTN passthrough mode.</li> <li>• <b>no-pass-through</b>—Do not enable OTN passthrough mode.</li> </ul>          | detail<br>extensive |
| OTN Line Loopback  | <p>Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: <b>enabled</b> or <b>disabled</b>.</p>                                                                                                                                                                                                                       | detail<br>extensive |
| OTN FEC statistics | <p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <ul style="list-style-type: none"> <li>• <b>Corrected Errors</b>—The count of corrected errors in the last second.</li> <li>• <b>Corrected Error Ratio</b>—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.</li> </ul> | detail<br>extensive |
| OTN FEC alarms     | <p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <ul style="list-style-type: none"> <li>• <b>FEC Degrade</b>—OTU FEC Degrade defect.</li> <li>• <b>FEC Excessive</b>—OTU FEC Excessive Error defect.</li> </ul>                                                                                                       | detail<br>extensive |
| OTN OC             | <p>OTN OC defects triggered on the interface.</p> <ul style="list-style-type: none"> <li>• <b>LOS</b>—OC Loss of Signal defect.</li> <li>• <b>LOF</b>—OC Loss of Frame defect.</li> <li>• <b>LOM</b>—OC Loss of Multiframe defect.</li> <li>• <b>Wavelength Lock</b>—OC Wavelength Lock defect.</li> </ul>                                         | detail<br>extensive |

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                             |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| <b>OTN OTU</b>          | <p>OTN OTU defects detected on the interface</p> <ul style="list-style-type: none"> <li>• <b>AIS</b>—OTN AIS alarm.</li> <li>• <b>BDI</b>—OTN OTU BDI alarm.</li> <li>• <b>IAE</b>—OTN OTU IAE alarm.</li> <li>• <b>TTIM</b>—OTN OTU TTIM alarm.</li> <li>• <b>SF</b>—OTN ODU bit error rate fault alarm.</li> <li>• <b>SD</b>—OTN ODU bit error rate defect alarm.</li> <li>• <b>TCA-ES</b>—OTN ODU ES threshold alarm.</li> <li>• <b>TCA-SES</b>—OTN ODU SES threshold alarm.</li> <li>• <b>TCA-UAS</b>—OTN ODU UAS threshold alarm.</li> <li>• <b>TCA-BBE</b>—OTN ODU BBE threshold alarm.</li> <li>• <b>BIP</b>—OTN ODU BIP threshold alarm.</li> <li>• <b>BBE</b>—OTN OTU BBE threshold alarm.</li> <li>• <b>ES</b>—OTN OTU ES threshold alarm.</li> <li>• <b>SES</b>—OTN OTU SES threshold alarm.</li> <li>• <b>UAS</b>—OTN OTU UAS threshold alarm.</li> </ul> | <b>detail<br/>extensive</b> |
| <b>Received DAPI</b>    | Destination Access Port Interface (DAPI) from which the packets were received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail<br/>extensive</b> |
| <b>Received SAPI</b>    | Source Access Port Interface (SAPI) from which the packets were received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail<br/>extensive</b> |
| <b>Transmitted DAPI</b> | Destination Access Port Interface (DAPI) to which the packets were transmitted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail<br/>extensive</b> |
| <b>Transmitted SAPI</b> | Source Access Port Interface (SAPI) to which the packets were transmitted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail<br/>extensive</b> |

**PCS statistics**

(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.

- **Bit errors**—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.
- **Errored blocks**—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.

**detail  
extensive**

## MAC statistics

**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following:

- **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.
- **Unicast packets**, **Broadcast packets**, and **Multicast packets**—Number of unicast, broadcast, and multicast packets.
- **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).
- **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.
- **MAC control frames**—Number of MAC control frames.
- **MAC pause frames**—Number of MAC control frames with **pause** operational code.
- **Oversized frames**—Number of frames that exceed 1518 octets.
- **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.
- **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.

extensive

|                                       |                                                                                                                                                                                                                                                                                                                                                                  |                  |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                       | <ul style="list-style-type: none"> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”</li> </ul> |                  |
| <b>OTN Received Overhead Bytes</b>    | APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08                                                                                                                                                                                                                                                                                | <b>extensive</b> |
| <b>OTN Transmitted Overhead Bytes</b> | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08                                                                                                                                                                                                                                                                                | <b>extensive</b> |

**Filter statistics**

**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.

- **Input packet count**—Number of packets received from the MAC hardware that the filter processed.
- **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address.
- **Input DA rejects**—Number of packets that the filter rejected because the destination 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 routing device from the far-end system, either there is a bad 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 routing device (which the routing device is rejecting).
- **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect.
- **Output packet count**—Number of packets that the filter has given to the MAC hardware.
- **Output packet pad count**—Number of packets 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 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.

**extensive**

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                  |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                | <ul style="list-style-type: none"> <li>• <b>Output packet error count</b>—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.</li> <li>• <b>CAM destination filters, CAM source filters</b>—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.</li> </ul> |                  |
| <b>PMA PHY</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. Any state other than <b>OK</b> indicates a problem.</li> </ul>                                                                                                                                                                                                                                                   | <b>extensive</b> |

**WIS section**

(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:

- **Seconds**—Number of seconds the defect has been active.
- **Count**—Number of times that the defect has gone from inactive to active.
- **State**—State of the error. Any state other than **OK** indicates a problem.

Subfields are:

- **BIP-B1**—Bit interleaved parity for SONET section overhead
- **SEF**—Severely errored framing
- **LOL**—Loss of light
- **LOF**—Loss of frame
- **ES-S**—Errored seconds (section)
- **SES-S**—Severely errored seconds (section)
- **SEFS-S**—Severely errored framing seconds (section)

**extensive**

**WIS line**

(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.

**extensive**

- **Seconds**—Number of seconds the defect has been active.
- **Count**—Number of times that the defect has gone from inactive to active.
- **State**—State of the error. State other than **OK** indicates a problem.

Subfields are:

- **BIP-B2**—Bit interleaved parity for SONET line overhead
- **REI-L**—Remote error indication (near-end line)
- **RDI-L**—Remote defect indication (near-end line)
- **AIS-L**—Alarm indication signal (near-end line)
- **BERR-SF**—Bit error rate fault (signal failure)
- **BERR-SD**—Bit error rate defect (signal degradation)
- **ES-L**—Errored seconds (near-end line)
- **SES-L**—Severely errored seconds (near-end line)
- **UAS-L**—Unavailable seconds (near-end line)
- **ES-LFE**—Errored seconds (far-end line)
- **SES-LFE**—Severely errored seconds (far-end line)
- **UAS-LFE**—Unavailable seconds (far-end line)

**WIS path**

(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.

**extensive**

- **Seconds**—Number of seconds the defect has been active.
- **Count**—Number of times that the defect has gone from inactive to active.
- **State**—State of the error. Any state other than **OK** indicates a problem.

Subfields are:

- **BIP-B3**—Bit interleaved parity for SONET section overhead
- **REI-P**—Remote error indication
- **LOP-P**—Loss of pointer (path)
- **AIS-P**—Path alarm indication signal
- **RDI-P**—Path remote defect indication
- **UNEQ-P**—Path unequipped
- **PLM-P**—Path payload label mismatch
- **ES-P**—Errored seconds (near-end STS path)
- **SES-P**—Severely errored seconds (near-end STS path)
- **UAS-P**—Unavailable seconds (near-end STS path)
- **SES-PFE**—Severely errored seconds (far-end STS path)
- **UAS-PFE**—Unavailable seconds (far-end STS path)

**Autonegotiation information**

Information about link autonegotiation.

**extensive**

- **Negotiation status:**
  - **Incomplete**—Ethernet interface has the speed or link mode configured.
  - **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
  - **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
- **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
- **Link partner:**
  - **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
  - **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
  - **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
- **Local resolution**—Information from the link partner:

|                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                  |
|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                                    | <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> |                  |
| <b>Received path trace, Transmitted path trace</b> | <p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>                                                                                     | <b>extensive</b> |
| <b>Packet Forwarding Engine configuration</b>      | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b> |

|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                  |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>CoS information</b> | <p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b> |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|

### Logical Interface

|                          |                                                                                    |                                           |
|--------------------------|------------------------------------------------------------------------------------|-------------------------------------------|
| <b>Logical interface</b> | Name of the logical interface.                                                     | All levels                                |
| <b>Index</b>             | Index number of the logical interface, which reflects its initialization sequence. | <b>detail</b><br><b>extensive</b><br>none |
| <b>SNMP ifIndex</b>      | SNMP interface index number for the logical interface.                             | <b>detail</b><br><b>extensive</b><br>none |
| <b>Generation</b>        | Unique number for use by Juniper Networks technical support only.                  | <b>detail</b><br><b>extensive</b>         |

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                 |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| <b>Flags</b>    | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels                                      |
| <b>VLAN-Tag</b> | <p>Rewrite profile applied to incoming or outgoing frames on the outer (<b>Out</b>) VLAN tag or for both the outer and inner (<b>In</b>) VLAN tags.</p> <ul style="list-style-type: none"> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>pop</b>—The outer VLAN tag of the incoming frame is removed.</li> <li>• <b>swap</b>—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information.</li> <li>• <b>push</b>—An outer VLAN tag is pushed in front of the existing VLAN tag.</li> <li>• <b>push-push</b>—Two VLAN tags are pushed in from the incoming frame.</li> <li>• <b>swap-push</b>—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.</li> <li>• <b>swap-swap</b>—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value.</li> <li>• <b>pop-swap</b>—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.</li> <li>• <b>pop-pop</b>—Both the outer and inner VLAN tags of the incoming frame are removed.</li> </ul> | <b>brief detail</b><br><b>extensive</b><br>none |

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                           |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <b>Demux:</b>                  | <p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> <li>• Source Family Inet</li> <li>• Destination Family Inet</li> </ul>                                                                                                                                                                                                                                                                                                                                | <b>detail</b><br><b>extensive</b><br>none |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels                                |
| <b>Protocol</b>                | Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b><br><b>extensive</b><br>none |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail</b><br><b>extensive</b><br>none |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail</b><br><b>extensive</b><br>none |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul> | <b>detail</b><br><b>extensive</b>         |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>extensive</b>                          |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| <b>Local statistics</b>         | Number and rate of bytes and packets destined to the routing device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b>                     |
| <b>Transit statistics</b>       | Number and rate of bytes and packets transiting the switch.<br><br><b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler. | <b>extensive</b>                     |
| <b>Generation</b>               | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail<br/>extensive</b>          |
| <b>Route Table</b>              | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail<br/>extensive<br/>none</b> |
| <b>Flags</b>                    | Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail<br/>extensive</b>          |
| <b>Donor interface</b>          | (Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail<br/>extensive<br/>none</b> |
| <b>Preferred source address</b> | (Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail<br/>extensive<br/>none</b> |
| <b>Input Filters</b>            | Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail<br/>extensive</b>          |

|                               |                                                                                                                                                                                                      |                                           |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------|
| <b>Output Filters</b>         | Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces. | <b>detail</b><br><b>extensive</b>         |
| <b>Mac-Validate Failures</b>  | Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.                                           | <b>detail</b><br><b>extensive</b><br>none |
| <b>Addresses, Flags</b>       | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                                                  | <b>detail</b><br><b>extensive</b><br>none |
| <b><i>protocol-family</i></b> | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                             | <b>brief</b>                              |
| <b>Flags</b>                  | Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> ).                                                      | <b>detail</b><br><b>extensive</b><br>none |
| <b>Destination</b>            | IP address of the remote side of the connection.                                                                                                                                                     | <b>detail</b><br><b>extensive</b><br>none |
| <b>Local</b>                  | IP address of the logical interface.                                                                                                                                                                 | <b>detail</b><br><b>extensive</b><br>none |
| <b>Broadcast</b>              | Broadcast address of the logical interlace.                                                                                                                                                          | <b>detail</b><br><b>extensive</b><br>none |
| <b>Generation</b>             | Unique number for use by Juniper Networks technical support only.                                                                                                                                    | <b>detail</b><br><b>extensive</b>         |

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. The following table describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500

bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). The **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit **.50** (VLAN 50).

**Table 155: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type**

| Interface Type             | Sample Command                               | Byte and Octet Counts Include                                                                                                                                                                                | Comments                                |
|----------------------------|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| Inbound physical interface | <b>show interfaces ge-0/3/0 extensive</b>    | Traffic statistics:<br><br>Input bytes: 496 bytes per packet, representing the Layer 2 packet<br><br>MAC statistics:<br><br>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes | The additional 4 bytes are for the CRC. |
| Inbound logical interface  | <b>show interfaces ge-0/3/0.50 extensive</b> | Traffic statistics:<br><br>Input bytes: 478 bytes per packet, representing the Layer 3 packet                                                                                                                |                                         |

**Table 155: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type**  
(Continued)

| Interface Type              | Sample Command                               | Byte and Octet Counts Include                                                                                                                                                                         | Comments                                                                                                                                           |
|-----------------------------|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Outbound physical interface | <b>show interfaces ge-0/0/0 extensive</b>    | Traffic statistics:<br>Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes<br><br>MAC statistics:<br>Received octets: 478 bytes per packet, representing the Layer 3 packet | For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type. |
| Outbound logical interface  | <b>show interfaces ge-0/0/0.50 extensive</b> | Traffic statistics:<br>Input bytes: 478 bytes per packet, representing the Layer 3 packet                                                                                                             |                                                                                                                                                    |

Table 156 on page 1524 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 156: show interfaces Output Fields**

| Field Name                | Field Description               | Level of Output |
|---------------------------|---------------------------------|-----------------|
| <b>Physical Interface</b> |                                 |                 |
| <b>Physical interface</b> | Name of the physical interface. | All levels      |
| <b>Enabled</b>            | State of the interface.         | All levels      |

Table 156: show interfaces Output Fields (Continued)

| Field Name              | Field Description                                                                                                              | Level of Output                 |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Interface index</b>  | Index number of the physical interface, which reflects its initialization sequence.                                            | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>     | SNMP index number for the physical interface.                                                                                  | <b>detail extensive</b><br>none |
| <b>Link-level type</b>  | Encapsulation being used on the physical interface.                                                                            | All levels                      |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                              | <b>detail extensive</b>         |
| <b>MTU</b>              | Maximum transmission unit size on the physical interface.                                                                      | All levels                      |
| <b>Link mode</b>        | Link mode: Full-duplex or Half-duplex.                                                                                         |                                 |
| <b>Speed</b>            | Speed at which the interface is running.                                                                                       | All levels                      |
| <b>BPDU error</b>       | Bridge protocol data unit (BPDU) error: Detected or None                                                                       |                                 |
| <b>Loopback</b>         | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> . | All levels                      |
| <b>Source filtering</b> | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .                                                                   | All levels                      |
| <b>Flow control</b>     | Flow control status: <b>Enabled</b> or <b>Disabled</b> .                                                                       | All levels                      |
| <b>Auto-negotiation</b> | (Gigabit Ethernet interfaces) Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .                                      | All levels                      |

Table 156: show interfaces Output Fields (*Continued*)

| Field Name             | Field Description                                                                                                                                                                                                                                           | Level of Output                 |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Remote-fault</b>    | (Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>         | All levels                      |
| <b>Device flags</b>    | Information about the physical device.                                                                                                                                                                                                                      | All levels                      |
| <b>Interface flags</b> | Information about the interface.                                                                                                                                                                                                                            | All levels                      |
| <b>Link flags</b>      | Information about the physical link.                                                                                                                                                                                                                        | All levels                      |
| <b>CoS queues</b>      | Number of CoS queues configured.                                                                                                                                                                                                                            | <b>detail extensive</b><br>none |
| <b>Current address</b> | Configured MAC address.                                                                                                                                                                                                                                     | <b>detail extensive</b><br>none |
| <b>Last flapped</b>    | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive</b><br>none |
| <b>Input Rate</b>      | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                           | None                            |
| <b>Output Rate</b>     | Output rate in bps and pps.                                                                                                                                                                                                                                 | None                            |

Table 156: show interfaces Output Fields (Continued)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output                 |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul> | <b>detail extensive</b><br>none |
| <b>Statistics last cleared</b>          | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>         |
| <b>Traffic statistics</b>               | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                  | <b>detail extensive</b>         |

Table 156: show interfaces Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output  |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b> | <p data-bbox="475 369 808 396">Input errors on the interface.</p> <ul style="list-style-type: none"> <li data-bbox="475 432 1154 499">• <b>Errors</b>—Sum of the incoming frame terminated and FCS errors.</li> <li data-bbox="475 537 1214 684">• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li data-bbox="475 722 1203 789">• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li data-bbox="475 827 1208 894">• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li data-bbox="475 932 1146 1079">• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li data-bbox="475 1117 1227 1306">• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <code>ignore-l3-incompletes</code> .</li> <li data-bbox="475 1344 1227 1411">• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li data-bbox="475 1449 1198 1558">• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li data-bbox="475 1596 1211 1705">• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li data-bbox="475 1743 980 1770">• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |

Table 156: show interfaces Output Fields (Continued)

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output  |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Output errors</b> | <p>Output errors on the interface.</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminated and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; therefore, for Gigabit Ethernet PICs, this number must always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces.</li> </ul> | <b>extensive</b> |

Table 156: show interfaces Output Fields (Continued)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                   | Level of Output         |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
|                                        | <ul style="list-style-type: none"> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>                                                                                                                                                      |                         |
| <b>Ingress queues</b>                  | Total number of ingress queues supported on the specified interface.                                                                                                                                                                                                                                                                                | <b>extensive</b>        |
| <b>Queue counters and queue number</b> | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> | <b>detail extensive</b> |

Table 156: show interfaces Output Fields (*Continued*)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output  |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>MAC statistics</b> | <p><b>Receive</b> and <b>Transmit</b> statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets.</li> <li>• <b>Unicast packets</b>, <b>Broadcast packets</b>, and <b>Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> <li>• Packet length exceeds 1518 octets, or</li> <li>• Packet length exceeds MRU</li> </ul> </li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> </ul> | <b>extensive</b> |

Table 156: show interfaces Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”</li> </ul> |                 |

Table 156: show interfaces Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output  |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Filter statistics</b> | <p><b>Receive</b> and <b>Transmit</b> statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> <li>• <b>Input packet count</b>—Number of packets received from the MAC hardware that the filter processed.</li> <li>• <b>Input packet rejects</b>—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address.</li> <li>• <b>Input DA rejects</b>—Number of packets that the filter rejected because the destination 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 device from the far-end system, either there is a bad 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 device (which the router is rejecting).</li> <li>• <b>Input SA rejects</b>—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect.</li> <li>• <b>Output packet count</b>—Number of packets that the filter has given to the MAC hardware.</li> <li>• <b>Output packet pad count</b>—Number of packets 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 IP packets</li> </ul> | <b>extensive</b> |

Table 156: show interfaces Output Fields (Continued)

| Field Name                                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output  |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                               | <p>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.</p> <ul style="list-style-type: none"> <li>• <b>Output packet error count</b>—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment.</li> <li>• <b>CAM destination filters, CAM source filters</b>—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0.</li> </ul> |                  |
| <b>Autonegotiation information</b>            | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> </ul>                                                                                                                                                                                                                                                                  | <b>extensive</b> |
| <b>Packet Forwarding Engine configuration</b> | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b> |

Table 156: show interfaces Output Fields (Continued)

| Field Name                           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>CoS information</b>               | <p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>        |
| <b>Interface transmit statistics</b> | Status of the <b>interface-transmit-statistics</b> configuration: Enabled or Disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <b>Queue counters (Egress)</b>       | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Logical Interface</b>             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                         |

Table 156: show interfaces Output Fields (Continued)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output                 |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Logical interface</b>  | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | All levels                      |
| <b>Index</b>              | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>       | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b><br>none |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>         |
| <b>Flags</b>              | Information about the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                      |
| <b>Encapsulation</b>      | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels                      |
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface set.</li> </ul> | <b>detail extensive</b>         |
| <b>Local statistics</b>   | Number and rate of bytes and packets destined to the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>extensive</b>                |

Table 156: show interfaces Output Fields (Continued)

| Field Name                                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                 |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Transit statistics</b>                             | <p>Number and rate of bytes and packets transiting the switch.</p> <p><b>NOTE:</b> For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the <b>Output bytes</b> and <b>Output packets</b> interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p> | <b>extensive</b>                |
| <b>Security</b>                                       | Security zones that interface belongs to.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b>                |
| <b>Flow Input statistics</b>                          | Statistics on packets received by flow module.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>extensive</b>                |
| <b>Flow Output statistics</b>                         | Statistics on packets sent by flow module.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>extensive</b>                |
| <b>Flow error statistics (Packets dropped due to)</b> | Statistics on errors in the flow module.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>extensive</b>                |
| <b>Protocol</b>                                       | Protocol family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b><br>none |
| <b>MTU</b>                                            | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b><br>none |
| <b>Generation</b>                                     | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |

Table 156: show interfaces Output Fields (Continued)

| Field Name              | Field Description                                                                                                 | Level of Output                 |
|-------------------------|-------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Route Table</b>      | Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0. | <b>detail extensive</b><br>none |
| <b>Flags</b>            | Information about protocol family flags. .                                                                        | <b>detail extensive</b>         |
| <b>Addresses, Flags</b> | Information about the address flags..                                                                             | <b>detail extensive</b><br>none |
| <b>Destination</b>      | IP address of the remote side of the connection.                                                                  | <b>detail extensive</b><br>none |
| <b>Local</b>            | IP address of the logical interface.                                                                              | <b>detail extensive</b><br>none |
| <b>Broadcast</b>        | Broadcast address of the logical interface.                                                                       | <b>detail extensive</b><br>none |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                 | <b>detail extensive</b>         |

### Sample output for G.fast and Annex J support

#### show interfaces (SRX380, SRX300, SRX320, SRX340, and SRX345)

```

user@host> show interfaces ge-0/0/8
Physical interface: ge-0/0/8, Enabled, Physical link is Up
 Interface index: 146, SNMP ifIndex: 520
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Full-duplex,
 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: Disabled, Auto-negotiation: Enabled, Remote fault: Online

```

## DSL SFP Status:

Chip Type : xDSL  
 Chip Firmware Version : 1\_62\_8463  
 Training Status : Showtime  
 Training Mode : ADSL2PLUS  
 Annex Type : Annex J  
 Profile Type : NA  
 Carrier Set : NA  
 Line Status : No Defect

| DSL SFP Statistics:      | XTU-R (DS) | XTU-C (US) |
|--------------------------|------------|------------|
| Packet Count :           | 0          | 0          |
| CRC Error Count :        | 0          | 0          |
| Electrical Length (dB) : | 0          | 0          |
| Net Data Rate (Kbps) :   | 25737      | 3143       |
| SNR Margin (dB) :        | 100        | -2         |
| CV Count :               | 0          | 0          |
| ES Count :               | 0          | 0          |
| SES Count :              | 0          | 0          |
| UAS Count :              | 0          | 0          |

Device flags : Present Running

Interface flags: SNMP-Traps Internal: 0x0

Link flags : None

CoS queues : 8 supported, 8 maximum usable queues

Current address: 4c:16:fc:de:30:89, Hardware address: 4c:16:fc:de:30:89

Last flapped : 2020-10-28 19:56:29 PDT (3d 23:07 ago)

Input rate : 20544 bps (42 pps)

Output rate : 20544 bps (42 pps)

Active alarms : None

Active defects : None

| PCS statistics | Seconds |
|----------------|---------|
|----------------|---------|

|            |   |
|------------|---|
| Bit errors | 0 |
|------------|---|

|                |   |
|----------------|---|
| Errored blocks | 0 |
|----------------|---|

| Ethernet FEC statistics | Errors |
|-------------------------|--------|
|-------------------------|--------|

|                      |   |
|----------------------|---|
| FEC Corrected Errors | 0 |
|----------------------|---|

|                        |   |
|------------------------|---|
| FEC Uncorrected Errors | 0 |
|------------------------|---|

|                           |   |
|---------------------------|---|
| FEC Corrected Errors Rate | 0 |
|---------------------------|---|

|                             |   |
|-----------------------------|---|
| FEC Uncorrected Errors Rate | 0 |
|-----------------------------|---|

Interface transmit statistics: Disabled

Logical interface ge-0/0/8.0 (Index 77) (SNMP ifIndex 538)

Flags: Up SNMP-Traps 0x0 VLAN-Tag [ 0x8100.10 ] Encapsulation: ENET2

Input packets : 0

Output packets: 105040

```

Security: Zone: trust
Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp sap vrrp dhcp finger ftp
tftp ident-reset http https ike netconf ping reverse-telnet reverse-ssh
rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
lsping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap
sdwan-appqoe l3-ha
Protocol inet, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 0,
Curr new hold cnt: 0, NH drop cnt: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 1.1.3/24, Local: 1.1.3.2, Broadcast: 1.1.3.255

```

### show interfaces (G.fast related information on SRX380, SRX300, SRX320, SRX340, and SRX345)

```

user@host> show interfaces ge-0/0/8
Physical interface: ge-0/0/8, Enabled, Physical link is Up
G.fast mode, DS Speed: 400Mbps, US Speed: 400Mbps
Cont.....
.....

```

### show interfaces terse (ACX5448, ACX5448-D, ACX710 channelized interface)

```

user@host> show interfaces terse et-0/1/2
Interface Admin Link Proto Local Remote
et-0/1/2:0 up down
et-0/1/2:1 up down
et-0/1/2:2 up down
et-0/1/2:3 up down

```

### show interfaces (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
Interface index: 167, SNMP ifIndex: 35
Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,

```

```

Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues : 4 supported, 4 maximum usable queues
Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
Last flapped : 2006-08-10 17:25:10 PDT (00:01:08 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
Active alarms : None
Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
 Flags: SNMP-Traps 0x4000
 VLAN-Tag [0x8100.512 0x8100.513] In(pop-swap 0x8100.530) Out(swap-push
 0x8100.512 0x8100.513)
 Encapsulation: VLAN-CCC
 Egress account overhead: 100
 Ingress account overhead: 90
 Input packets : 0
 Output packets: 0
 Protocol ccc, MTU: 1522
 Flags: Is-Primary

```

## show interfaces (Gigabit Ethernet on MX Series Routers)

```

user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
 Interface index: 156, SNMP ifIndex: 188
 Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, 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, 4 maximum usable queues
Schedulers : 0
Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76

```

```

Last flapped : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Active alarms : None
Active defects : None
Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
 Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
 Input packets : 10232
 Output packets: 10294
 Protocol inet, MTU: 1500
 Flags: Sendbcast-pkt-to-re
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast:
203.0.113.255 Protocol inet6, MTU: 1500
 Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold
cnt: 4, NH drop cnt: 0
 Flags: Is-Primary
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 2001:db8:/32, Local: 2001:db8::5
 Addresses, Flags: Is-Preferred
 Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
 Protocol multiservice, MTU: Unlimited
 Flags: Is-Primary

```

### show interfaces (link degrade status)

```

user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
 Interface index: 157, SNMP ifIndex: 537
 Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU Error:
None, Loopback: Disabled, 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: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
 Last flapped : 2014-06-18 02:36:38 PDT (02:50:50 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 0
Link Degraded* :
Link Monitoring : Enable
Link Degraded Set Threshold: : 1E-7
Link Degraded Clear Threshold: : 1E-12
Estimated BER : 1E-7
Link-degraded event : Seconds Count State
 782 1 Defect Active

```

### show interfaces extensive (link degrade status) (PTX10001-36MR)

```

user@host> show interfaces et-0/0/1 extensive
Physical interface: et-0/0/1, Enabled, Physical link is Down
 Interface index: 1017, SNMP ifIndex: 519, Generation: 712964572820
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU
Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None, Loopback:
Disabled,
 Source filtering: Disabled, Flow control: Enabled, Media type: Fiber
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x80
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Damping : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
state: unsuppressed
Current address: 40:de:ad:28:7a:0a, Hardware address: 40:de:ad:28:7a:0a
Last flapped : 2020-08-27 12:05:18 IST (00:50:56 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 2239638274139000 392152080440 bps
 Output bytes : 0 0 bps
 Input packets: 2239638274049 49019008 pps
 Output packets: 0 0 pps
Input errors:
 Errors: 1, 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: 2, 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 0 0 0
 1 0 0 0
 2 0 0 0
 3 0 0 0
Queue number: Mapped forwarding classes
 0 best-effort
 1 expedited-forwarding
 2 assured-forwarding
 3 network-control
Active alarms : LINK
Active defects : LINK, LOCAL-DEGRADE
PCS statistics Seconds
 Bit errors 0
 Errored blocks 0
Ethernet FEC Mode : FEC119
Ethernet FEC statistics Errors
 FEC Corrected Errors 166615427699
 FEC Uncorrected Errors 12
 FEC Corrected Errors Rate 3687323
 FEC Uncorrected Errors Rate 0
MAC statistics: Receive Transmit
 Total octets 2239648398609518 0
 Total packets 2239648398691 0
 Unicast packets 2239648398036 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0
 VLAN tagged frames 0
 Code violations 0
 Total errors 1 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0

```

```

Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
 Destination slot: ()
CoS information:
 Direction :
Interface transmit statistics: Disabled
Link Degradate :
 Link Monitoring : Enable
 Link Degradate Set Threshold : 1E-5
 Link Degradate Clear Threshold : 1E-10
 Link Degradate War Set Threshold : 1E-9
 Link Degradate War Clear Threshold : 1E-11
 Estimated BER : 1E-5
 Link-degrade event : Seconds
Count State
1 Defect Active
3054

```

### show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```

user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
Interface index: 151, SNMP ifIndex: 530, Generation: 154
Interface flags: SNMP-Traps Internal: 0x4000
Output bytes : 240614363944 772721536 bps
Output packets: 35384446506 1420444 pps
Direction : Output
Interface transmit statistics: Enabled

Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
Output bytes : 195560312716 522726272 bps
Output packets: 4251311146 1420451 pps

user@host> show interfaces ge-5/2/0.0 statistics detail
Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2

```

```

Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
 Input bytes : 271524
 Output bytes : 37769598
 Input packets: 3664
 Output packets: 885790
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 16681118
 Input packets: 0
 Output packets: 362633
Local statistics:
 Input bytes : 271524
 Output bytes : 308560
 Input packets: 3664
 Output packets: 3659
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 37461038 0 bps
 Input packets: 0 0 pps
 Output packets: 882131 0 pps
IPv6 transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 16681118 0 bps
 Input packets: 0 0 pps
 Output packets: 362633 0 pps

```

### show interfaces brief (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
 Link-level type: 52, MTU: 1522, Speed: 1000mbps, 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

Logical interface ge-3/0/2.0
 Flags: SNMP-Traps 0x4000

```

```
VLAN-Tag [0x8100.512 0x8100.513] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
Encapsulation: VLAN-CCC
ccc
```

```
Logical interface ge-3/0/2.32767
```

```
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x0000.0] Encapsulation: ENET2
```

## show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
```

```
Physical interface: ge-3/0/2, Enabled, Physical link is Up
```

```
Interface index: 167, SNMP ifIndex: 35, Generation: 177
```

```
Link-level type: 52, MTU: 1522, Speed: 1000mbps, 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 : 4 supported, 4 maximum usable queues
```

```
Hold-times : Up 0 ms, Down 0 ms
```

```
Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
```

```
Last flapped : 2006-08-09 17:17:00 PDT (01:31:33 ago)
```

```
Statistics last cleared: Never
```

```
Traffic statistics:
```

```
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
```

```
Ingress traffic statistics at Packet Forwarding Engine:
```

```
Input bytes : 0 0 bps
Input packets: 0 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
```

```
Ingress queues: 4 supported, 4 in use
```

| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
|-----------------|----------------|---------------------|-----------------|
| 0 best-effort   | 0              | 0                   | 0               |
| 1 expedited-fo  | 0              | 0                   | 0               |
| 2 assured-forw  | 0              | 0                   | 0               |
| 3 network-cont  | 0              | 0                   | 0               |

```
Egress queues: 4 supported, 4 in use
```

| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
|-----------------|----------------|---------------------|-----------------|
| 0 best-effort   | 0              | 0                   | 0               |
| 1 expedited-fo  | 0              | 0                   | 0               |
| 2 assured-forw  | 0              | 0                   | 0               |
| 3 network-cont  | 0              | 0                   | 0               |

Active alarms : None

Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)

Flags: SNMP-Traps 0x4000

VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push 0x8100.512 0x8100.513)

Encapsulation: VLAN-CCC

Egress account overhead: 100

Ingress account overhead: 90

Traffic statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Local statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Transit statistics:

|                 |   |       |
|-----------------|---|-------|
| Input bytes :   | 0 | 0 bps |
| Output bytes :  | 0 | 0 bps |
| Input packets:  | 0 | 0 pps |
| Output packets: | 0 | 0 pps |

Protocol ccc, MTU: 1522, Generation: 149, Route table: 0

Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)

(Generation 139)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

Traffic statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Local statistics:

|               |   |
|---------------|---|
| Input bytes : | 0 |
|---------------|---|

```

Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

```

### show interfaces extensive (Gigabit Ethernet IQ2)

```

user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
Last flapped : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 38910844056 7952 bps
Output bytes : 7174605 8464 bps
Input packets: 418398473 11 pps
Output packets: 78903 12 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 38910799145 7952 bps
Input packets: 418397956 11 pps
Drop bytes : 0 0 bps
Drop packets: 0 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: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,  
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0

Ingress queues: 4 supported, 4 in use

| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
|-----------------|----------------|---------------------|-----------------|
| 0 best-effort   | 418390823      | 418390823           | 0               |
| 1 expedited-fo  | 0              | 0                   | 0               |
| 2 assured-forw  | 0              | 0                   | 0               |
| 3 network-cont  | 7133           | 7133                | 0               |

Egress queues: 4 supported, 4 in use

| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
|-----------------|----------------|---------------------|-----------------|
| 0 best-effort   | 1031           | 1031                | 0               |
| 1 expedited-fo  | 0              | 0                   | 0               |
| 2 assured-forw  | 0              | 0                   | 0               |
| 3 network-cont  | 77872          | 77872               | 0               |

Active alarms : None

Active defects : None

| MAC statistics:                                                | Receive      | Transmit                     |
|----------------------------------------------------------------|--------------|------------------------------|
| Total octets                                                   | 38910844056  | 7174605                      |
| Total packets                                                  | 418398473    | 78903                        |
| Unicast packets                                                | 408021893366 | 1026                         |
| Broadcast packets                                              | 10           | 12                           |
| Multicast packets                                              | 418398217    | 77865                        |
| 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            | OTN Received Overhead Bytes: |
| APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 |              |                              |
| Payload Type: 0x08                                             |              |                              |

OTN Transmitted Overhead Bytes:

APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00  
 Payload Type: 0x08

Filter statistics:

Input packet count 418398473

```

Input packet rejects 479
Input DA rejects 479
Input SA rejects 0
Output packet count 78903
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: 7
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
 0 best-effort 95 950000000 95 0
low none
 3 network-control 5 500000000 5 0
low none
 Direction : Input
 CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
 0 best-effort 95 950000000 95 0
low none
 3 network-control 5 500000000 5 0
low none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
 Input bytes : 812400
 Output bytes : 1349206
 Input packets: 9429
 Output packets: 9449
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0

```

```

Output packets: 0
Local statistics:
Input bytes : 812400
Output bytes : 1349206
Input packets: 9429
Output packets: 9449
Transit statistics:
Input bytes : 0 7440 bps
Output bytes : 0 7888 bps
Input packets: 0 10 pps
Output packets: 0 11 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
 Flags: Is-Primary, Mac-Validate-Strict
 Mac-Validate Failures: Packets: 0, Bytes: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
 Output Filters: F2-ge-3/0/1.0-out (53)
 Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
 Generation: 196
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
 Flags: Is-Primary
 Policer: Input: __default_arp_policer__

```

**NOTE:** For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in [Table 154 on page 1468](#).

### show interfaces (Gigabit Ethernet Unnumbered Interface)

```

user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
Interface index: 148, SNMP ifIndex: 50

```

```

Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, 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, 4 maximum usable queues
Current address: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53:f8
Last flapped : 2006-10-27 04:42:23 PDT (08:01:52 ago)
Input rate : 0 bps (0 pps)
Output rate : 624 bps (1 pps)
Active alarms : None
Active defects : None

```

```

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
 Flags: SNMP-Traps Encapsulation: ENET2
 Input packets : 0
 Output packets: 6
 Protocol inet, MTU: 1500
 Flags: Unnumbered
 Donor interface: lo0.0 (Index 64)
 Preferred source address: 203.0.113.22

```

## show interfaces (ACI Interface Set Configured)

```

user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
 Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.4001] Encapsulation: PPP-over-
Ethernet
ACI VLAN:
 Dynamic Profile: aci-vlan-set-profile
 PPPoE:
 Dynamic Profile: aci-vlan-pppoe-profile,
 Service Name Table: None,
 Max Sessions: 32000, Max Sessions VSA Ignore: Off,
 Duplicate Protection: On, Short Cycle Protection: Off,
 Direct Connect: Off,
 AC Name: nbc
 Input packets : 9

```

```
Output packets: 8
Protocol multiservice, MTU: Unlimited
```

## show interfaces (ALI Interface Set)

```
user@host> show interfaces ge-1/0/0.10
Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [0x8100.10] Encapsulation: ENET2
Line Identity:
 Dynamic Profile: ali-set-profile
 Circuit-id Remote-id Accept-no-ids
PPPoE:
 Dynamic Profile: ali-vlan-pppoe-profile,
 Service Name Table: None,
 Max Sessions: 32000, Max Sessions VSA Ignore: Off,
 Duplicate Protection: On, Short Cycle Protection: Off,
 Direct Connect: Off,
 AC Name: nbc
Input packets : 9
Output packets: 8
Protocol multiservice, MTU: Unlimited
```

## Sample Output Gigabit Ethernet

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```
user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
Interface index: 177, SNMP ifIndex: 630, Generation: 178
Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Enabled,
Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 1024
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53:f6
```

```

Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 6970332384 0 bps
 Output bytes : 0 0 bps
 Input packets: 81050506 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
 Input bytes : 6970299398 0 bps
 Input packets: 81049992 0 pps
 Drop bytes : 0 0 bps
 Drop packets: 0 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
Ingress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
 0 best-effort 81049992 81049992 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont 0 0 0
Egress queues: 4 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
 0 best-effort 0 0 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont 0 0 0
Active alarms : None
Active defects : None
PCS statistics Seconds
 Bit errors 0
 Errored blocks 0
MAC statistics: Receive Transmit

```

```

Total octets 6970332384 0
Total packets 81050506 0
Unicast packets 81050000 0
Broadcast packets 506 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0

Filter statistics:
Input packet count 81050506
Input packet rejects 506
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Packet Forwarding Engine configuration:
Destination slot: 5

CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none
Direction : Input
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % usec
0 best-effort 95 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.100] Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
Input bytes : 0

```

```

Output bytes : 46
Input packets: 0
Output packets: 1
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 46
 Input packets: 0
 Output packets: 1
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 253, Route table: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 265
 Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__

```

### show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 630, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 9.294GbpsGbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled Speed Configuration: Auto
Device flags : Present Running
Interface flags: SNMP-Traps 16384

```

```

Link flags : None
CoS queues : 4 supported
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
Last flapped : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps
 Output packets: 0 0 pps
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,
 HS Link CRC errors: 0, HS Link FIFO overflows: 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
Queue counters:
 Queued packets Transmitted packets Dropped packets
 0 best-effort 0 0 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont 0 0 0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics
 Seconds Count
 Bit errors 0 0
 Errored blocks 0 0
MAC statistics:
 Receive Transmit
 Total octets 0 0
 Total packets 0 0
 Unicast packets 0 0
 Broadcast packets 0 0
 Multicast packets 0 0
 CRC/Align errors 0 0
 FIFO errors 0 0
 MAC control frames 0 0
 MAC pause frames 0 0
 Oversized frames 0
 Jabber frames 0
 Fragment frames 0

```

```

VLAN tagged frames 0
Code violations 0
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
PMA PHY:
 Seconds Count State
 PLL lock 0 0 OK
 PHY light 63159 1 Light Missing
WIS section:
 BIP-B1 0 0
 SEF 434430 434438 Defect Active
 LOS 434430 1 Defect Active
 LOF 434430 1 Defect Active
 ES-S 434430
 SES-S 434430
 SEFS-S 434430
WIS line:
 BIP-B2 0 0
 REI-L 0 0
 RDI-L 0 0 OK
 AIS-L 434430 1 Defect Active
 BERR-SF 0 0 OK
 BERR-SD 0 0 OK
 ES-L 434430
 SES-L 434430
 UAS-L 434420
 ES-LFE 0
 SES-LFE 0
 UAS-LFE 0
WIS path:
 BIP-B3 0 0
 REI-P 0 0
 LOP-P 0 0 OK
 AIS-P 434430 1 Defect Active
 RDI-P 0 0 OK
 UNEQ-P 0 0 OK
 PLM-P 0 0 OK

```

```

ES-P 434430
SES-P 434430
UAS-P 434420
ES-PFE 0
SES-PFE 0
UAS-PFE 0

```

Received path trace:

```
00
```

Transmitted path trace: orissa so-1/0/0

```
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
```

Packet Forwarding Engine configuration:

Destination slot: 1

CoS information:

| CoS | transmit queue  | Bandwidth | Buffer    | Priority | Limit      |
|-----|-----------------|-----------|-----------|----------|------------|
|     | %               | bps       | % bytes   |          |            |
| 0   | best-effort     | 95        | 950000000 | 95       | 0 low none |
| 3   | network-control | 5         | 50000000  | 5        | 0 low none |

## show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```
user@host> show interfaces ge-7/0/0 extensive
```

Physical interface: ge-7/0/0, Enabled, Physical link is Down

Interface index: 143, SNMP ifIndex: 508, Generation: 208

Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,

MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,

Flow control: Enabled

Device flags : Present Running Down

Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000

Link flags : None

Wavelength : 1550.12 nm, Frequency: 193.40 THz

CoS queues : 8 supported, 8 maximum usable queues

Hold-times : Up 0 ms, Down 0 ms

Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72

Last flapped : 2011-04-20 15:48:54 PDT (18:39:49 ago)

Statistics last cleared: Never

Traffic statistics:

```

Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

```

IPv6 transit statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
 FIFO errors: 0, Resource errors: 0
Output errors:
 Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
 0 best-effort 0 0 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont
Queue number: Mapped forwarding classes
 0 best-effort
 1 expedited-forwarding
 2 assured-forwarding
 3 network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
 Receive Transmit
Total octets 0 0
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Total octets 0 0
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0

```

```

Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0

OTN alarms : None
OTN defects : None
OTN FEC Mode : GFEC
OTN Rate : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
 Corrected Errors 0
 Corrected Error Ratio (0 sec average) 0e-0
OTN FEC alarms:
 Seconds Count State
 FEC Degrade 0 0 OK
 FEC Excessive 0 0 OK
OTN OC:
 Seconds Count State
 LOS 2 1 OK
 LOF 67164 2 Defect Active
 LOM 67164 71 Defect Active
 Wavelength Lock 0 0 OK
OTN OTU:
 AIS 0 0 OK
 BDI 65919 4814 Defect Active
 IAE 67158 1 Defect Active
 TTIM 7 1 OK
 SF 67164 2 Defect Active
 SD 67164 3 Defect Active
 TCA-ES 0 0 OK
 TCA-SES 0 0 OK
 TCA-UAS 80 40 OK
 TCA-BBE 0 0 OK
 BIP 0 0 OK
 BBE 0 0 OK
 ES 0 0 OK
 SES 0 0 OK
 UAS 587 0 OK
Received DAPI:

```

```

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Received SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
OTN Received Overhead Bytes:
 APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
 Payload Type: 0x03
OTN Transmitted Overhead Bytes:
 APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
 Payload Type: 0x03
Filter statistics:
 Input packet count 0
 Input packet rejects 0
 Input DA rejects 0
 Input SA rejects 0
 Output packet count 0
 Output packet pad count 0
 Output packet error count 0
 CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
 Destination slot: 7
CoS information:
 Direction : Output
 CoS transmit queue Bandwidth Buffer Priority
Limit
 % bps % usec
 0 best-effort 95 9500000000 95 0 low
none
 3 network-control 5 500000000 5 0 low
none
 ...

```

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```

user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
Interface index: 173, SNMP ifIndex: 212, Generation: 174
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,

```

```

Unidirectional: Enabled,
 Loopback: None, Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
...

```

### show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
 Interface index: 176, SNMP ifIndex: 137, Generation: 177
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Tx-Only
 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: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
 Last flapped : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 322891152287160 9627472888 bps
 Input packets: 0 0 pps
 Output packets: 328809727380 1225492 pps
...

 Filter statistics:
 Output packet count 328810554250
 Output packet pad count 0
 Output packet error count 0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)
 Flags: SNMP-Traps Encapsulation: ENET2
 Egress account overhead: 100
 Ingress account overhead: 90
 Traffic statistics:
 Input bytes : 0

```

```

Output bytes : 322891152287160
Input packets: 0
Output packets: 328809727380
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 322891152287160 9627472888 bps
 Input packets: 0 0 pps
 Output packets: 328809727380 1225492 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__

```

### **show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)**

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
 Interface index: 174, SNMP ifIndex: 118, Generation: 175
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
 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: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
Last flapped : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 322857456303482 9627496104 bps
 Output bytes : 0 0 bps
 Input packets: 328775413751 1225495 pps
 Output packets: 0 0 pps

```

...

```

Filter statistics:
 Input packet count 328775015056
 Input packet rejects 1
 Input DA rejects 0

```

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

```

 Input bytes : 322857456303482
 Output bytes : 0
 Input packets: 328775413751
 Output packets: 0

```

IPv6 transit statistics:

```

 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

Local statistics:

```

 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

Transit statistics:

```

 Input bytes : 322857456303482 9627496104 bps
 Output bytes : 0 0 bps
 Input packets: 328775413751 1225495 pps
 Output packets: 0 0 pps

```

```

IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
Protocol inet, MTU: 1500, Generation: 145, Route table: 0
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 139
Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
 Flags: None
 Policer: Input: __default_arp_policer__

```

## Sample Output

### Sample Output SRX Gigabit Ethernet

```

user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
 Interface index: 135, SNMP ifIndex: 510
 Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
 Last flapped : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : LINK
 Active defects: LINK
 Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Security: Zone: public

```

```

Protocol inet, MTU: 1500
 Flags: Sendbroadcast-pkt-to-re
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255

```

## Sample Output SRX Gigabit Ethernet

```

user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
 Interface index: 135, SNMP ifIndex: 510
 Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
 Last flapped : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : LINK
 Active defects: LINK
 Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
 Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Security: Zone: public
 Protocol inet, MTU: 1500
 Flags: Sendbroadcast-pkt-to-re
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255

```

**show interfaces (Gigabit Ethernet for vSRX and vSRX 3.0)**

```

user@host> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Up
 Interface index: 136, SNMP ifIndex: 510
 Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Link-mode: Half-duplex,
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
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 CoS queues : 8 supported, 8 maximum usable queues
 Current address: 00:50:56:93:ef:25, Hardware address: 00:50:56:93:ef:25
 Last flapped : 2019-03-29 01:57:45 UTC (00:00:41 ago)
 Input rate : 1120 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Active alarms : None

```

**show interfaces detail (Gigabit Ethernet)**

```

user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
 Interface index: 135, SNMP ifIndex: 510, Generation: 138
 Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
 Last flapped : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps

```

```

Output packets: 0 0 pps
Egress queues: 8 supported, 4 in use
Queue counters: Queued packets Transmitted packets Dropped packets
 0 best-effort 0 0 0
 1 expedited-fo 0 0 0
 2 assured-forw 0 0 0
 3 network-cont 0 0 0
Queue number: Mapped forwarding classes
 0 best-effort
 1 expedited-forwarding
 2 assured-forwarding
 3 network-control
Active alarms : LINK
Active defects : LINK
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
 Self packets : 0
 ICMP packets : 0
 VPN packets : 0
 Multicast packets : 0
 Bytes permitted by policy : 0
 Connections established : 0
Flow Output statistics:

```

```

Multicast packets : 0
Bytes permitted by policy : 0
Flow error statistics (Packets dropped due to):
Address spoofing: 0
Authentication failed: 0
Incoming NAT errors: 0
Invalid zone received packet: 0
Multiple user authentications: 0
Multiple incoming NAT: 0
No parent for a gate: 0
No one interested in self packets: 0
No minor session: 0
No more sessions: 0
No NAT gate: 0
No route present: 0
No SA for incoming SPI: 0
No tunnel found: 0
No session for a gate: 0
No zone or NULL zone binding 0
Policy denied: 0
Security association not active: 0
TCP sequence number out of window: 0
Syn-attack protection: 0
User authentication errors: 0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
 Flags: Sendbroadcast-pkt-to-re
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150

```

## show interfaces statistics st0.0 detail

```

user@host> show interfaces statistics st0.0 detail
Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
 Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
 Traffic statistics:
 Input bytes : 528152756774
 Output bytes : 575950643520
 Input packets: 11481581669
 Output packets: 12520666095
 Local statistics:

```

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 121859888 bps
Output bytes : 0 128104112 bps
Input packets: 0 331141 pps
Output packets: 0 348108 pps
Security: Zone: untrust
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
sap vrrp
Flow Statistics :
Flow Input statistics :
Self packets : 0
ICMP packets : 0
VPN packets : 0
Multicast packets : 0
Bytes permitted by policy : 525984295844
Connections established : 7
Flow Output statistics:
Multicast packets : 0
Bytes permitted by policy : 576003290222
Flow error statistics (Packets dropped due to):
Address spoofing: 0
Authentication failed: 0
Incoming NAT errors: 0
Invalid zone received packet: 0
Multiple user authentications: 0
Multiple incoming NAT: 0
No parent for a gate: 0
No one interested in self packets: 0
No minor session: 0
No more sessions: 0
No NAT gate: 0
No route present: 2000280
No SA for incoming SPI: 0
No tunnel found: 0
No session for a gate: 0
No zone or NULL zone binding 0
Policy denied: 0
Security association not active: 0

```

```

 TCP sequence number out of window: 0
 Syn-attack protection: 0
 User authentication errors: 0
 Protocol inet, MTU: 9192
 Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
 NH drop cnt: 0
 Generation: 155, Route table: 0
 Flags: Sendbcast-pkt-to-re

```

### show interfaces extensive (Gigabit Ethernet)

```

user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
 Interface index: 135, SNMP ifIndex: 510, Generation: 138
 Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, 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 Down
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
 Last flapped : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets : 0 0 pps
 Output packets: 0 0 pps
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
 L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
 FIFO errors: 0, Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
 Egress queues: 8 supported, 4 in use
 Queue counters: Queued packets Transmitted packets Dropped packets

```

```

0 best-effort 0 0 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0
Queue number: Mapped forwarding classes
0 best-effort
1 expedited-forwarding
2 assured-forwarding
3 network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
Total octets Receive Transmit
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0
CRC/Align errors 0 0
FIFO errors 0 0
MAC control frames 0 0
MAC pause frames 0 0
Oversized frames 0
Jabber frames 0
Fragment frames 0
VLAN tagged frames 0
Code violations 0
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 2, CAM source filters: 0
Autonegotiation information:
Negotiation status: Incomplete
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue Bandwidth Buffer Priority
Limit

```

|                   | %  | bps       | %  | usec |     |
|-------------------|----|-----------|----|------|-----|
| 0 best-effort     | 95 | 950000000 | 95 | 0    | low |
| none              |    |           |    |      |     |
| 3 network-control | 5  | 50000000  | 5  | 0    | low |
| none              |    |           |    |      |     |

Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

Input bytes : 0  
Output bytes : 0  
Input packets: 0  
Output packets: 0

Local statistics:

Input bytes : 0  
Output bytes : 0  
Input packets: 0  
Output packets: 0

Transit statistics:

Input bytes : 0 0 bps  
Output bytes : 0 0 bps  
Input packets: 0 0 pps  
Output packets: 0 0 pps

Security: Zone: public

Flow Statistics :

Flow Input statistics :

Self packets : 0  
ICMP packets : 0  
VPN packets : 0  
Multicast packets : 0  
Bytes permitted by policy : 0  
Connections established : 0

Flow Output statistics:

Multicast packets : 0  
Bytes permitted by policy : 0

Flow error statistics (Packets dropped due to):

Address spoofing: 0  
Authentication failed: 0  
Incoming NAT errors: 0  
Invalid zone received packet: 0  
Multiple user authentications: 0  
Multiple incoming NAT: 0

```

No parent for a gate: 0
No one interested in self packets: 0
No minor session: 0
No more sessions: 0
No NAT gate: 0
No route present: 0
No SA for incoming SPI: 0
No tunnel found: 0
No session for a gate: 0
No zone or NULL zone binding 0
Policy denied: 0
Security association not active: 0
TCP sequence number out of window: 0
Syn-attack protection: 0
User authentication errors: 0

Protocol inet, MTU: 1500, Generation: 150, Route table: 0
 Flags: Sendbcst-pkt-to-re
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
 Generation: 150

```

### show interfaces terse

```

user@host> show interfaces terse

Interface Admin Link Proto Local Remote
ge-0/0/0 up up
ge-0/0/0.0 up up inet 10.209.4.61/18
gr-0/0/0 up up
ip-0/0/0 up up
st0 up up
st0.1 up ready inet
ls-0/0/0 up up
lt-0/0/0 up up
mt-0/0/0 up up
pd-0/0/0 up up
pe-0/0/0 up up
e3-1/0/0 up up
t3-2/0/0 up up
e1-3/0/0 up up

```

```

se-4/0/0 up down
t1-5/0/0 up up
br-6/0/0 up up
dc-6/0/0 up up
dc-6/0/0.32767 up up
bc-6/0/0:1 down up
bc-6/0/0:1.0 up down
dl0 up up
dl0.0 up up inet
dsc up up
gre up up
ipip up up
lo0 up up
lo0.16385 up up inet 10.0.0.1 --> 0/0
 10.0.0.16 --> 0/0

lsi up up
mtun up up
pimd up up
pime up up
pp0 up up

```

### show interfaces terse (vSRX and vSRX 3.0)

```

user@host> show interfaces terse
Interface Admin Link Proto Local Remote
ge-0/0/0 up up
ge-0/0/0.0 up up inet 1.1.65.1/24
ge-0/0/1 up up
ge-0/0/2 up up
e-0/0/3 up up
ge-0/0/4 up up

```

### show interfaces controller (Channelized E1 IQ with Logical E1)

```

user@host> show interfaces controller ce1-1/2/6

Controller Admin Link
ce1-1/2/6 up up
e1-1/2/6 up up

```

**show interfaces controller (Channelized E1 IQ with Logical DS0)**

```
user@host> show interfaces controller ce1-1/2/3
```

| Controller | Admin | Link |
|------------|-------|------|
| ce1-1/2/3  | up    | up   |
| ds-1/2/3:1 | up    | up   |
| ds-1/2/3:2 | up    | up   |

**show interfaces descriptions**

```
user@host> show interfaces descriptions
```

| Interface | Admin | Link | Description      |
|-----------|-------|------|------------------|
| so-1/0/0  | up    | up   | M20-3#1          |
| so-2/0/0  | up    | up   | GSR-12#1         |
| ge-3/0/0  | up    | up   | SMB-OSPF_Area300 |
| so-3/3/0  | up    | up   | GSR-13#1         |
| so-3/3/1  | up    | up   | GSR-13#2         |
| ge-4/0/0  | up    | up   | T320-7#1         |
| ge-5/0/0  | up    | up   | T320-7#2         |
| so-7/1/0  | up    | up   | M160-6#1         |
| ge-8/0/0  | up    | up   | T320-7#3         |
| ge-9/0/0  | up    | up   | T320-7#4         |
| so-10/0/0 | up    | up   | M160-6#2         |
| so-13/0/0 | up    | up   | M20-3#2          |
| so-14/0/0 | up    | up   | GSR-12#2         |
| ge-15/0/0 | up    | up   | SMB-OSPF_Area100 |
| ge-15/0/1 | up    | up   | GSR-13#3         |

**show interfaces destination-class all**

```
user@host> show interfaces destination-class all
```

```
Logical interface so-4/0/0.0
```

| Destination class | Packets<br>(packet-per-second) | Bytes<br>(bits-per-second) |
|-------------------|--------------------------------|----------------------------|
| gold              | 0                              | 0                          |
|                   | ( 0) (                         | 0)                         |
| silver            | 0                              | 0                          |
|                   | ( 0) (                         | 0)                         |

```

Logical interface so-0/1/3.0

 Packets Bytes
Destination class (packet-per-second) (bits-per-second)
 gold 0 0
 (0) (0)
 silver 0 0
 (0) (0)

```

### show interfaces diagnostics optics

```

user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
 Laser bias current : 7.408 mA
 Laser output power : 0.3500 mW / -4.56 dBm
 Module temperature : 23 degrees C / 73 degrees F
 Module voltage : 3.3450 V
 Receiver signal average optical power : 0.0002 mW / -36.99 dBm
 Laser bias current high alarm : Off
 Laser bias current low alarm : Off
 Laser bias current high warning : Off
 Laser bias current low warning : Off
 Laser output power high alarm : Off
 Laser output power low alarm : Off
 Laser output power high warning : Off
 Laser output power low warning : Off
 Module temperature high alarm : Off
 Module temperature low alarm : Off
 Module temperature high warning : Off
 Module temperature low warning : Off
 Module voltage high alarm : Off
 Module voltage low alarm : Off
 Module voltage high warning : Off
 Module voltage low warning : Off
 Laser rx power high alarm : Off
 Laser rx power low alarm : On
 Laser rx power high warning : Off
 Laser rx power low warning : On
 Laser bias current high alarm threshold : 17.000 mA
 Laser bias current low alarm threshold : 1.000 mA
 Laser bias current high warning threshold : 14.000 mA
 Laser bias current low warning threshold : 2.000 mA

```

```

Laser output power high alarm threshold : 0.6310 mW / -2.00 dBm
Laser output power low alarm threshold : 0.0670 mW / -11.74 dBm
Laser output power high warning threshold : 0.6310 mW / -2.00 dBm
Laser output power low warning threshold : 0.0790 mW / -11.02 dBm
Module temperature high alarm threshold : 95 degrees C / 203 degrees F
Module temperature low alarm threshold : -25 degrees C / -13 degrees F
Module temperature high warning threshold : 90 degrees C / 194 degrees F
Module temperature low warning threshold : -20 degrees C / -4 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7940 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

### show interfaces far-end-interval coc12-5/2/0

```

user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
05:30-current:
 ES-L: 1, SES-L: 1, UAS-L: 0
05:15-05:30:
 ES-L: 0, SES-L: 0, UAS-L: 0
05:00-05:15:
 ES-L: 0, SES-L: 0, UAS-L: 0
04:45-05:00:
 ES-L: 0, SES-L: 0, UAS-L: 0
04:30-04:45:
 ES-L: 0, SES-L: 0, UAS-L: 0
04:15-04:30:
 ES-L: 0, SES-L: 0, UAS-L: 0
04:00-04:15:
...

```

## show interfaces far-end-interval coc1-5/2/1:1

```

user@host> run show interfaces far-end-interval coc1-5/2/1:1
Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
05:30-current:
 ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:15-05:30:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:00-05:15:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:45-05:00:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:30-04:45:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:15-04:30:
 ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:00-04:15:

```

## show interfaces filters

```

user@host> show interfaces filters
Interface Admin Link Proto Input Filter Output Filter
ge-0/0/0 up up
ge-0/0/0.0 up up inet
 iso
ge-5/0/0 up up
ge-5/0/0.0 up up any f-any
 inet f-inet
 multiservice
gr-0/3/0 up up
ip-0/3/0 up up
mt-0/3/0 up up
pd-0/3/0 up up
pe-0/3/0 up up
vt-0/3/0 up up
at-1/0/0 up up
at-1/0/0.0 up up inet
 iso
at-1/1/0 up down
at-1/1/0.0 up down inet

```



```

No SA for incoming SPI: 0
No tunnel found: 0
No session for a gate: 0
No zone or NULL zone binding 0
Policy denied: 0
Security association not active: 0
TCP sequence number out of window: 0
Syn-attack protection: 0
User authentication errors: 0
Protocol inet, MTU: 1500
 Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255

```

### show interfaces interval (Channelized OC12)

```

user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
17:43-current:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
17:28-17:43:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
17:13-17:28:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
16:58-17:13:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
16:43-16:58:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
 LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
 CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

**show interfaces interval (E3)**

```

user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
 17:43-current:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 17:28-17:43:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 17:13-17:28:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 16:58-17:13:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 16:43-16:58:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,

Interval Total:
 LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
 CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

**show interfaces interval (SONET/SDH) (SRX devices)**

```

user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
 20:02-current:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
 SES-P: 0, UAS-P: 0
 19:47-20:02:
 ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
 ES-P: 267, SES-P: 267, UAS-P: 267
 19:32-19:47:
 ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
 SES-P: 56, UAS-P: 46
 19:17-19:32:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
 SES-P: 0, UAS-P: 0

```

```
19:02-19:17:
.....
```

### show interfaces load-balancing (SRX devices)

```
user@host> show interfaces load-balancing
Interface State Last change Member count
ams0 Up 1d 00:50 2
ams1 Up 00:00:59 2
```

### show interfaces load-balancing detail (SRX devices)

```
user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface : ams0
State : Up
Last change : 1d 00:51
Member count : 2
Members :
 Interface Weight State
 mams-2/0/0 10 Active
 mams-2/1/0 10 Active
```

### show interfaces mac-database (All MAC Addresses on a Port SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
 Interface index: 372, SNMP ifIndex: 788
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
```

| MAC address       | Input frames | Input bytes | Output frames | Output bytes |
|-------------------|--------------|-------------|---------------|--------------|
| 00:00:00:00:00:00 | 1            | 56          | 0             | 0            |
| 00:00:c0:01:01:02 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:03 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:04 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:05 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:06 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:07 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:08 | 7023809      | 323095214   | 0             | 0            |
| 00:00:c0:01:01:09 | 7023809      | 323095214   | 0             | 0            |
| 00:00:c0:01:01:0a | 7023809      | 323095214   | 0             | 0            |
| 00:00:c0:01:01:0b | 7023809      | 323095214   | 0             | 0            |
| 00:00:c8:01:01:02 | 30424784     | 1399540064  | 37448598      | 1722635508   |
| 00:00:c8:01:01:03 | 30424784     | 1399540064  | 37448598      | 1722635508   |
| 00:00:c8:01:01:04 | 30424716     | 1399536936  | 37448523      | 1722632058   |
| 00:00:c8:01:01:05 | 30424789     | 1399540294  | 37448598      | 1722635508   |
| 00:00:c8:01:01:06 | 30424788     | 1399540248  | 37448597      | 1722635462   |
| 00:00:c8:01:01:07 | 30424783     | 1399540018  | 37448597      | 1722635462   |
| 00:00:c8:01:01:08 | 30424783     | 1399540018  | 37448596      | 1722635416   |
| 00:00:c8:01:01:09 | 8836796      | 406492616   | 8836795       | 406492570    |
| 00:00:c8:01:01:0a | 30424712     | 1399536752  | 37448521      | 1722631966   |
| 00:00:c8:01:01:0b | 30424715     | 1399536890  | 37448523      | 1722632058   |

Number of MAC addresses : 21

### show interfaces mac-database (All MAC Addresses on a Service SRX devices)

```
user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
```

| MAC address       | Input frames | Input bytes | Output frames | Output bytes |
|-------------------|--------------|-------------|---------------|--------------|
| 00:00:00:00:00:00 | 1            | 56          | 0             | 0            |
| 00:00:c0:01:01:02 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:03 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:04 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:05 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:06 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:07 | 7023810      | 323095260   | 0             | 0            |
| 00:00:c0:01:01:08 | 7023809      | 323095214   | 0             | 0            |
| 00:00:c0:01:01:09 | 7023809      | 323095214   | 0             | 0            |
| 00:00:c0:01:01:0a | 7023809      | 323095214   | 0             | 0            |

|                   |          |            |          |            |
|-------------------|----------|------------|----------|------------|
| 00:00:c0:01:01:0b | 7023809  | 323095214  | 0        | 0          |
| 00:00:c8:01:01:02 | 31016568 | 1426762128 | 38040381 | 1749857526 |
| 00:00:c8:01:01:03 | 31016568 | 1426762128 | 38040382 | 1749857572 |
| 00:00:c8:01:01:04 | 31016499 | 1426758954 | 38040306 | 1749854076 |
| 00:00:c8:01:01:05 | 31016573 | 1426762358 | 38040381 | 1749857526 |
| 00:00:c8:01:01:06 | 31016573 | 1426762358 | 38040381 | 1749857526 |
| 00:00:c8:01:01:07 | 31016567 | 1426762082 | 38040380 | 1749857480 |
| 00:00:c8:01:01:08 | 31016567 | 1426762082 | 38040379 | 1749857434 |
| 00:00:c8:01:01:09 | 9428580  | 433714680  | 9428580  | 433714680  |
| 00:00:c8:01:01:0a | 31016496 | 1426758816 | 38040304 | 1749853984 |
| 00:00:c8:01:01:0b | 31016498 | 1426758908 | 38040307 | 1749854122 |

### show interfaces mac-database mac-address

```

user@host> show interfaces mac-database xe-0/3/3 mac-address (SRX devices)
00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
 Interface index: 372, SNMP ifIndex: 788
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address: 00:00:c8:01:01:09, Type: Configured,
 Input bytes : 202324652
 Output bytes : 202324560
 Input frames : 4398362
 Output frames : 4398360
Policer statistics:
Policer type Discarded frames Discarded bytes
Output aggregate 3992386 183649756

```

## show interfaces mc-ae (SRX devices)

```

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

```

## show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```

user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
Interface index: 168, SNMP ifIndex: 495
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps 16384
Link flags : Keepalives
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
LCP state: Opened
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Not-configured
CoS queues : 8 supported
Last flapped : 2005-06-15 12:14:59 PDT (04:31:29 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
SONET alarms : None
SONET defects : None
SONET errors:
 BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
Received path trace: routerb so-1/1/2
Transmitted path trace: routera so-4/1/2

```

**show interfaces policers (SRX devices)**

```

user@host> show interfaces policers
Interface Admin Link Proto Input Policer Output Policer
ge-0/0/0 up up
ge-0/0/0.0 up up inet
 iso
gr-0/3/0 up up
ip-0/3/0 up up
mt-0/3/0 up up
pd-0/3/0 up up
pe-0/3/0 up up
...
so-2/0/0 up up
so-2/0/0.0 up up inet so-2/0/0.0-in-policer so-2/0/0.0-out-policer
 iso
so-2/1/0 up down
...

```

**show interfaces policers interface-name (SRX devices)**

```

user@host> show interfaces policers so-2/1/0
Interface Admin Link Proto Input Policer Output Policer
so-2/1/0 up down
so-2/1/0.0 up down inet so-2/1/0.0-in-policer so-2/1/0.0-out-policer
 iso
 inet6

```

**show interfaces queue (SRX devices)**

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```

user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
 Interface index: 134, SNMP ifIndex: 509
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: class0

```

```

Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RL-dropped packets : 0 0 pps
 RL-dropped bytes : 0 0 bps
 RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue Buffer Usage:
 Reserved buffer : 118750000 bytes
Queue-depth bytes :
 Current : 0
..
..
Queue: 1, Forwarding classes: class1
..
..
Queue Buffer Usage:
 Reserved buffer : 9192 bytes
Queue-depth bytes :
 Current : 0
..
..
Queue: 3, Forwarding classes: class3
 Queued:
 ..
 ..
Queue Buffer Usage:
 Reserved buffer : 6250000 bytes
Queue-depth bytes :
 Current : 0

```

```
..
..
```

### show interfaces redundancy (SRX devices)

```
user@host> show interfaces redundancy
Interface State Last change Primary Secondary Current status
rsp0 Not present sp-1/0/0 sp-0/2/0 both down
rsp1 On secondary 1d 23:56 sp-1/2/0 sp-0/3/0 primary down
rsp2 On primary 10:10:27 sp-1/3/0 sp-0/2/0 secondary down
rlsq0 On primary 00:06:24 lsq-0/3/0 lsq-1/0/0 both up
```

### show interfaces redundancy (Aggregated Ethernet SRX devices)

```
user@host> show interfaces redundancy
Interface State Last change Primary Secondary Current status
rlsq0 On secondary 00:56:12 lsq-4/0/0 lsq-3/0/0 both up
ae0
ae1
ae2
ae3
ae4
```

### show interfaces redundancy detail (SRX devices)

```
user@host> show interfaces redundancy detail
Interface : rlsq0
State : On primary
Last change : 00:45:47
Primary : lsq-0/2/0
Secondary : lsq-1/2/0
Current status : both up
Mode : hot-standby

Interface : rlsq0:0
State : On primary
Last change : 00:45:46
Primary : lsq-0/2/0:0
```

```

Secondary : lsq-1/2/0:0
Current status : both up
Mode : warm-standby

```

### show interfaces routing brief (SRX devices)

```

user@host> show interfaces routing brief
Interface State Addresses
so-5/0/3.0 Down ISO enabled
so-5/0/2.0 Up MPLS enabled
 ISO enabled
 INET 192.168.2.120
 INET enabled
so-5/0/1.0 Up MPLS enabled
 ISO enabled
 INET 192.168.2.130
 INET enabled
at-1/0/0.3 Up CCC enabled
at-1/0/0.2 Up CCC enabled
at-1/0/0.0 Up ISO enabled
 INET 192.168.90.10
 INET enabled
lo0.0 Up ISO 47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
 ISO enabled
 INET 127.0.0.1
fxp1.0 Up
fxp0.0 Up INET 192.168.6.90

```

### show interfaces routing detail (SRX devices)

```

user@host> show interfaces routing detail
so-5/0/3.0
 Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>
 Metric: 0, Up/down transitions: 0, Full-duplex
 Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
 ISO address (null)
 State: <Broadcast PointToPoint Multicast> Change: <>
 Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0

```

```

Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>
Metric: 0, Up/down transitions: 0, Full-duplex
Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
MPLS address (null)
 State: <Up Broadcast PointToPoint Multicast> Change: <>
 Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
ISO address (null)
 State: <Up Broadcast PointToPoint Multicast> Change: <>
 Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
INET address 192.168.2.120
 State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
 Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
 Local address: 192.168.2.120
 Destination: 192.168.2.110/32
INET address (null)
 State: <Up Broadcast PointToPoint Multicast> Change: <>
 Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...

```

### show interfaces routing-instance all (SRX devices)

```

user@host> show interfaces terse routing-instance all
Interface Admin Link Proto Local Remote Instance
at-0/0/1 up up inet 10.0.0.1/24
ge-0/0/0.0 up up inet 192.168.4.28/24 sample-a
at-0/1/0.0 up up inet6 fe80::a:0:0:4/64 sample-b
so-0/0/0.0 up up inet 10.0.0.1/32

```

### show interfaces snmp-index (SRX devices)

```

user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
Interface index: 149, SNMP ifIndex: 33
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags : Present Running Down
Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
Link flags : Keepalives
CoS queues : 8 supported

```

```

Last flapped : 2005-06-15 11:45:57 PDT (05:38:43 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
SONET alarms : LOL, PLL, LOS
SONET defects : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P

```

### show interfaces source-class all (SRX devices)

```

user@host> show interfaces source-class all
Logical interface so-0/1/0.0

```

| Source class | Packets<br>(packet-per-second) | Bytes<br>(bits-per-second) |
|--------------|--------------------------------|----------------------------|
| gold         | 1928095                        | 161959980                  |
| (            | 889)                           | ( 597762)                  |
| bronze       | 0                              | 0                          |
| (            | 0)                             | ( 0)                       |
| silver       | 0                              | 0                          |
| (            | 0)                             | ( 0)                       |

```

Logical interface so-0/1/3.0

```

| Source class | Packets<br>(packet-per-second) | Bytes<br>(bits-per-second) |
|--------------|--------------------------------|----------------------------|
| gold         | 0                              | 0                          |
| (            | 0)                             | ( 0)                       |
| bronze       | 0                              | 0                          |
| (            | 0)                             | ( 0)                       |
| silver       | 116113                         | 9753492                    |
| (            | 939)                           | ( 631616)                  |

### show interfaces statistics (Fast Ethernet SRX devices)

```

user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
Interface index: 144, SNMP ifIndex: 1042
Description: ford fe-1/3/1
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues : 4 supported, 4 maximum usable queues

```

```

Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
Last flapped : 2006-04-18 03:08:59 PDT (00:01:24 ago)
Statistics last cleared: Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Input errors: 0, Output errors: 0
Active alarms : None
Active defects : None
Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500
 Flags: Is-Primary, DCU, SCU-in

 Packets Bytes
Destination class (packet-per-second) (bits-per-second)
 silver1 0 0
 (0) (0)
 silver2 0 0
 (0) (0)
 silver3 0 0
 (0) (0)

 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 10.27.245/24, Local: 10.27.245.2,
 Broadcast: 10.27.245.255
 Protocol iso, MTU: 1497
 Flags: Is-Primary

```

### show interfaces switch-port (SRX devices)

```

user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
 Speed: 100mbps, Auto-negotiation: Enabled
 Statistics:
 Total bytes 28437086 Transmit 21792250
 Total packets 409145 88008
 Unicast packets 9987 83817
 Multicast packets 145002 0
 Broadcast packets 254156 4191
 Multiple collisions 23 10
 FIFO/CRC/Align errors 0 0
 MAC pause frames 0 0
 Oversized frames 0

```

```

Runt frames 0
Jabber frames 0
Fragment frames 0
Discarded frames 0

Autonegotiation information:
Negotiation status: Complete
Link partner:
 Link mode: Full-duplex, Flow control: None, Remote fault: OK,
Link partner Speed: 100 Mbps
Local resolution:
 Flow control: None, Remote fault: Link OK

```

### show interfaces transport pm (SRX devices)

```

user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
 14:45-current Elapse time:900 Seconds
Near End Suspect Flag:False Reason:None
 PM COUNT THRESHOLD TCA-ENABLED TCA-
RAISED
 OTU-BBE 0 800 No No
 OTU-ES 0 135 No No
 OTU-SES 0 90 No No
 OTU-UAS 427 90 No No
Far End Suspect Flag:True Reason:Unknown
 PM COUNT THRESHOLD TCA-ENABLED TCA-
RAISED
 OTU-BBE 0 800 No No
 OTU-ES 0 135 No No
 OTU-SES 0 90 No No
 OTU-UAS 0 90 No No
Near End Suspect Flag:False Reason:None
 PM COUNT THRESHOLD TCA-ENABLED TCA-
RAISED
 ODU-BBE 0 800 No No
 ODU-ES 0 135 No No
 ODU-SES 0 90 No No
 ODU-UAS 427 90 No No
Far End Suspect Flag:True Reason:Unknown
 PM COUNT THRESHOLD TCA-ENABLED TCA-
RAISED

```

|                                                |                    |                        |             |           |
|------------------------------------------------|--------------------|------------------------|-------------|-----------|
| ODU-BBE                                        | 0                  | 800                    | No          | No        |
| ODU-ES                                         | 0                  | 135                    | No          | No        |
| ODU-SES                                        | 0                  | 90                     | No          | No        |
| ODU-UAS                                        | 0                  | 90                     | No          | No        |
| FEC                                            | Suspect Flag:False |                        | Reason:None |           |
| PM                                             | COUNT              | THRESHOLD              | TCA-ENABLED | TCA-      |
| RAISED                                         |                    |                        |             |           |
| FEC-CorrectedErr                               | 2008544300         | 0                      | NA          | NA        |
| FEC-UncorrectedWords                           | 0                  | 0                      | NA          | NA        |
| BER                                            | Suspect Flag:False |                        | Reason:None |           |
| PM                                             | MIN                | MAX                    | AVG         | THRESHOLD |
| ENABLED                                        | TCA-RAISED         |                        |             |           |
| BER                                            | 3.6e-5             | 5.8e-5                 | 3.6e-5      | 10.0e-3   |
| No                                             | Yes                |                        |             |           |
| Physical interface: et-0/1/0, SNMP ifIndex 515 |                    |                        |             |           |
| 14:45-current                                  |                    |                        |             |           |
| Suspect Flag:True                              |                    | Reason:Object Disabled |             |           |
| PM                                             | CURRENT            | MIN                    | MAX         | AVG       |
| THRESHOLD                                      | TCA-ENABLED        | TCA-RAISED             |             |           |
|                                                |                    |                        |             | (MIN)     |
| (MAX)                                          | (MIN)              | (MAX)                  | (MIN)       | (MAX)     |
| Lane chromatic dispersion                      | 0                  | 0                      | 0           | 0         |
| 0                                              | NA                 | NA                     | NA          | NA        |
| Lane differential group delay                  | 0                  | 0                      | 0           | 0         |
| 0                                              | NA                 | NA                     | NA          | NA        |
| q Value                                        |                    | 120                    | 120         | 120       |
| 0                                              | 0                  | NA                     | NA          | NA        |
| SNR                                            |                    | 28                     | 28          | 29        |
| 0                                              | NA                 | NA                     | NA          | NA        |
| Tx output power(0.01dBm)                       |                    | -5000                  | -5000       | -5000     |
| -100                                           | No                 | No                     | No          | No        |
| Rx input power(0.01dBm)                        |                    | -3642                  | -3665       | -3626     |
| -500                                           | No                 | No                     | No          | No        |
| Module temperature(Celsius)                    |                    | 46                     | 46          | 46        |
| 75                                             | No                 | No                     | No          | No        |
| Tx laser bias current(0.1mA)                   |                    | 0                      | 0           | 0         |
| 0                                              | NA                 | NA                     | NA          | NA        |
| Rx laser bias current(0.1mA)                   |                    | 1270                   | 1270        | 1270      |
| 0                                              | NA                 | NA                     | NA          | NA        |
| Carrier frequency offset(MHz)                  |                    | -186                   | -186        | -186      |
| 5000                                           | No                 | No                     | No          | No        |

## show security zones (SRX devices)

```
user@host> show security zones
Functional zone: management
 Description: This is the management zone.
 Policy configurable: No
 Interfaces bound: 1
 Interfaces:
 ge-0/0/0.0
Security zone: Host
 Description: This is the host zone.
 Send reset for non-SYN session TCP packets: Off
 Policy configurable: Yes
 Interfaces bound: 1
 Interfaces:
 fxp0.0
Security zone: abc
 Description: This is the abc zone.
 Send reset for non-SYN session TCP packets: Off
 Policy configurable: Yes
 Interfaces bound: 1
 Interfaces:
 ge-0/0/1.0
Security zone: def
 Description: This is the def zone.
 Send reset for non-SYN session TCP packets: Off
 Policy configurable: Yes
 Interfaces bound: 1
 Interfaces:
 ge-0/0/2.0
```

## show interfaces extensive (QFX5130-32CD)

```
user@host> show interfaces et-0/0/29 extensive | no-more
Physical interface: et-0/0/29, Enabled, Physical link is Up
 Interface index: 1086, SNMP ifIndex: 549, Generation: 618475300929
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps,
 BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
 Media type: Copper
```

```

Device flags : Present Running
Interface flags: SNMP-Traps
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Damping : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
state: unsuppressed
Current address: 0c:59:9c:81:fb:12, Hardware address: 0c:59:9c:81:fb:12
Last flapped : 2020-09-14 06:27:45 PDT (1d 01:34 ago)
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 130245061850190 0 bps
 Output bytes : 132765627331264 0 bps
 Input packets : 86830042098 0 pps
 Output packets: 88510419103 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: 4, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 5 in use
Queue counters:
 Queued packets Transmitted packets Dropped packets
 0 0 0 0
 3 0 0 0
 4 0 0 0
 7 88510418098 88510418098 0
 8 0 0 0
Queue number:
 Mapped forwarding classes
 0 best-effort
 3 fcoe
 4 no-loss
 7 network-control
 8 mcast
Active alarms : None
Active defects : None
PCS statistics
 Bit errors Seconds
 Bit errors 0
 Errored blocks 0
Ethernet FEC Mode : FEC119
Ethernet FEC statistics
 Errors
 FEC Corrected Errors 5796771110
 FEC Uncorrected Errors 0

```

```

FEC Corrected Errors Rate 0
FEC Uncorrected Errors Rate 0
MAC statistics:
 Receive Transmit
Total octets 130245061850190 132765627331264
Total packets 86830042098 88510419103
Unicast packets 86830041265 88510418256
Broadcast packets 179 190
Multicast packets 654 657
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
MAC Priority Flow Control Statistics:
Priority : 0 0 0
Priority : 1 0 0
Priority : 2 0 0
Priority : 3 0 0
Priority : 4 0 0
Priority : 5 0 0
Priority : 6 0 0
Priority : 7 0 0
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: ()
CoS information:
Direction :
Interface transmit statistics: Disabled

```

```
Link Degradе :
Link Monitoring : Disable
```

## Release Information

Command introduced before Junos OS Release 7.4.

Command modified in Junos OS Release 9.5 for SRX Series devices.

Command modified in Junos OS Release 19.3R1 for MX Series Routers.

Starting in Junos OS Release 19.3R1, Output fields **lindex** and **speed** is modified in the **show interfaces interface name extensive** command, on all MX Series routers.

- The default behavior of WAN-PHY interface remains the same. The new **precise-bandwidth** option reflects the new speed (9.294-Gbps) configured on the supported line cards.
- The WAN-PHY framing mode is supported only on MPC5E and MPC6E line cards.

Starting in Junos OS Release 19.3R1, class of service (CoS) features can be configured on the physical interface with speed rates of 1-Gbps, 10-Gbps, 40-Gbps, and 100-Gbps to provide better bandwidth for processing traffic during congestion using variant speeds.

## Release History Table

| Release | Description                                                                                                                                                                                                                                    |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 20.4R1  | Starting in Junos OS Release 20.4R1, we support G.fast and Annex J specification with SFP xDSL for ADSL2/ADSL2+ and all VDSL2 profiles on SRX380, SRX300, SRX320, SRX340, and SRX345 devices.                                                  |
| 19.2R3  | In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the show interfaces interface-name<media><extensive> command displays the autonegotiation status only for the interface that supports autonegotiation. |
| 18.4R1  | Starting in Junos OS Release 18.4R1, Output fields Next-hop and vpls-status is displayed in the show interfaces interface name detail command, only for Layer 2 protocols on MX480 routers.                                                    |

## RELATED DOCUMENTATION

[Understanding Layer 2 Interfaces on Security Devices](#)

[Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration](#)

[Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers](#)

[dsl-sfp-options](#)

## show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)

### IN THIS SECTION

- [Syntax \(M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface\) | 1602](#)
- [Syntax \(M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface\) | 1602](#)
- [Description | 1603](#)
- [Options | 1603](#)
- [Required Privilege Level | 1603](#)
- [Output Fields | 1603](#)
- [Sample Output | 1609](#)
- [Release Information | 1627](#)

### Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface)

```
show interfaces em0 | fxp0 | mgmtre0
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface)

```
show interfaces bcm0 | em0 | em1| fxp1 | fxp2 | ixgbe0 | ixgbe1
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Description

(M Series, T Series, TX Matrix Plus, and PTX Series devices only) Display status information about the management Ethernet and internal Ethernet interfaces.

## Options

**em0 | fxp0 | mgmtre0** (M Series, MX Series, T Series, and PTX Series) Display standard information about the management Ethernet interface. For supported Ethernet interface by chassis and Routing Engine, see [Supported Routing Engines by Router](#).

**bcm0 | em0 | em1 | fxp1 | fxp2 | ixgbe0 | ixgbe1** (M Series, MX Series, T Series, and PTX Series) Display standard information about the internal Ethernet interfaces. See [Supported Routing Engines by Router](#) for the internal Ethernet interface names for each Routing Engine by hardware platform.

**NOTE:** On Junos OS Evolved, the ixgbe0 and ixgbe1 internal interfaces are deprecated.

**brief | detail | extensive | terse** (Optional) Display the specified level of output.

**descriptions** (Optional) Display interface description strings.

**media** (Optional) Display media-specific information.

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

**statistics** (Optional) Display static interface statistics.

## Required Privilege Level

view

## Output Fields

[Table 157 on page 1604](#) lists the output fields for the **show interfaces** (management) command on the M Series routers, T Series routers, TX Matrix Plus routers, and PTX Series. Output fields are listed in the approximate order in which they appear.

**Table 157: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface**

| Field Name                | Field Description                                                                                                                    | Level of Output                 |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Physical Interface</b> |                                                                                                                                      |                                 |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                      | All levels                      |
| <b>Enabled</b>            | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels                      |
| <b>Interface index</b>    | Physical interface index number, which reflects its initialization sequence.                                                         | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                        | <b>detail extensive</b><br>none |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                    | <b>detail extensive</b>         |
| <b>Type</b>               | Type of interface.                                                                                                                   | All levels                      |
| <b>Link-level type</b>    | Encapsulation type used on the physical interface.                                                                                   | All levels                      |
| <b>MTU</b>                | Maximum transmission unit (MTU)—Size of the largest packet to be transmitted.                                                        | All levels                      |
| <b>Clocking</b>           | Reference clock source of the interface.                                                                                             | All levels                      |
| <b>Speed</b>              | Network speed on the interface.                                                                                                      | All levels                      |

**Table 157: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (Continued)**

| Field Name                    | Field Description                                                                                                                                  | Level of Output                 |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Device flags</b>           | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels                      |
| <b>Interface flags</b>        | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .    | All levels                      |
| <b>Link type</b>              | Data transmission type.                                                                                                                            | <b>detail extensive</b><br>none |
| <b>Link flags</b>             | Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .              | <b>detail extensive</b>         |
| <b>Physical info</b>          | Information about the physical interface.                                                                                                          | <b>detail extensive</b>         |
| <b>Hold-times</b>             | Current interface hold-time up and hold-time down. Value is in milliseconds.                                                                       | <b>detail extensive</b>         |
| <b>Current address</b>        | Configured MAC address.                                                                                                                            | <b>detail extensive</b><br>none |
| <b>Hardware address</b>       | Media access control (MAC) address of the interface.                                                                                               | <b>detail extensive</b><br>none |
| <b>Alternate link address</b> | Backup link address.                                                                                                                               | <b>detail extensive</b>         |

**Table 157: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (Continued)**

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                 | Level of Output                 |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                       | <b>detail extensive</b><br>none |
| <b>Input packets</b>           | Number of packets received on the physical interface.                                                                                                                                                                                                                                                                                                             | None specified                  |
| <b>Output packets</b>          | Number of packets transmitted on the physical interface.                                                                                                                                                                                                                                                                                                          | None specified                  |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b>         |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the logical and physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |

Table 157: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output  |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b>      | <ul style="list-style-type: none"> <li>• <b>Errors</b>—Input errors on the interface.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Frames received smaller than the runt threshold.</li> <li>• <b>Giants</b>—Frames received larger than the giant threshold.</li> <li>• <b>Policed Discards</b>—Frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>                                                                                                        | <b>extensive</b> |
| <b>Output errors</b>     | <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning.</li> <li>• <b>Errors</b>—Sum of outgoing frame terminates and FCS errors.</li> <li>• <b>Drops</b>—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 dropped by the ASIC RED mechanism.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                  |

**Table 157: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (Continued)**

| Field Name               | Field Description                                                                                                                           | Level of Output                 |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Logical interface</b> | Name of the logical interface                                                                                                               | All levels                      |
| <b>Index</b>             | Logical interface index number, which reflects its initialization sequence.                                                                 | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>      | Logical interface SNMP interface index number.                                                                                              | <b>detail extensive</b><br>none |
| <b>Generation</b>        | Unique number for use by Juniper Networks technical support only.                                                                           | <b>detail extensive</b>         |
| <b>Flags</b>             | Information about the logical interface; values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | All levels                      |
| <b>Encapsulation</b>     | Encapsulation on the logical interface.                                                                                                     | <b>detail extensive</b><br>none |
| <b>inet</b>              | IP address of the logical interface.                                                                                                        | <b>brief</b>                    |
| <b>Protocol</b>          | Protocol family configured on the logical interface (such as <b>iso</b> or <b>inet6</b> ).                                                  | <b>detail extensive</b><br>none |
| <b>MTU</b>               | MTU size on the logical interface.                                                                                                          | <b>detail extensive</b><br>none |
| <b>Generation</b>        | Unique number for use by Juniper Networks technical support only.                                                                           | <b>detail extensive</b>         |

**Table 157: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface (Continued)**

| Field Name              | Field Description                                                                                                                                        | Level of Output                 |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Route table</b>      | Route table in which this address exists. For example, <b>Route table:0</b> refers to inet.0.                                                            | <b>detail extensive</b>         |
| <b>Flags</b>            | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | <b>detail extensive</b><br>none |
| <b>Addresses, Flags</b> | Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .          | <b>detail extensive</b><br>none |
| <b>Destination</b>      | IP address of the remote side of the connection.                                                                                                         | <b>detail extensive</b><br>none |
| <b>Local</b>            | IP address of the logical interface.                                                                                                                     | <b>detail extensive</b><br>none |
| <b>Broadcast</b>        | Broadcast address.                                                                                                                                       | <b>detail extensive</b><br>none |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                        | <b>detail extensive</b>         |

## Sample Output

### show interfaces brief (Management Ethernet)

```

user@host> show interfaces fxp0 brief
Physical interface: fxp0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 100mbps

```

```

Device flags : Present Running
Interface flags: SNMP-Traps

Logical interface fxp0.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet 192.168.70.143/21

```

### show interfaces (Management Ethernet)

```

user@host> show interfaces fxp0
Physical interface: fxp0, Enabled, Physical link is Up
 Interface index: 1, SNMP ifIndex: 1
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Half-Duplex
 Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
 Last flapped : Never
 Input packets : 80804
 Output packets: 1105

Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.168.64/21, Local: 192.168.70.143,
 Broadcast: 192.168.71.255

```

### show interfaces (Management Ethernet [TX Matrix Plus Router])

```

user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
 Interface index: 8, SNMP ifIndex: 17
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0

```

```

Last flapped : Never
 Input packets : 1424
 Output packets: 5282

Logical interface em0.0 (Index 3) (SNMP ifIndex 18)
 Flags: SNMP-Traps Encapsulation: ENET2
 Input packets : 1424
 Output packets: 5282
 Protocol inet, MTU: 1500
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127

```

### show interfaces (Management Ethernet [PTX Series Packet Transport Routers])

```

user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
 Interface index: 8, SNMP ifIndex: 0
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
 Last flapped : Never
 Input packets : 212581
 Output packets: 71

 Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
 Flags: SNMP-Traps Encapsulation: ENET2
 Input packets : 212551
 Output packets: 71
 Protocol inet, MTU: 1500
 Flags: Is-Primary
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 192.168.3/24, Local: 192.168.3.30,
 Broadcast: 192.168.3.255

```

**show interfaces detail (Management Ethernet)**

```

user@host> show interfaces fxp0 detail
Physical interface: fxp0, Enabled, Physical link is Up
 Interface index: 1, SNMP ifIndex: 1, Generation: 0
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Half-Duplex
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 6484031
 Output bytes : 167503
 Input packets: 81008
 Output packets: 1110

Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500, Generation: 6, Route table: 0
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.168.64/21, Local: 192.168.70.143,
 Broadcast: 192.168.71.255, Generation: 1

```

**show interfaces detail (Management Ethernet [TX Matrix Plus Router])**

```

user@host> show interfaces em0 detail
Physical interface: em0, Enabled, Physical link is Up
 Interface index: 8, SNMP ifIndex: 17, Generation: 2
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex

```

```

Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
Alternate link address: Unspecified
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 124351
Output bytes : 1353212
Input packets: 1804
Output packets: 5344
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes : 117135
Output bytes : 1331647
Input packets: 1804
Output packets: 5344
Local statistics:
Input bytes : 117135
Output bytes : 1331647
Input packets: 1804
Output packets: 5344
Protocol inet, MTU: 1500, Generation: 1, Route table: 0
Flags: Is-Primary
Addresses, Flags: Is-Preferred Is-Primary
Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127, Generation: 1

```

### show interfaces detail (Management Ethernet [PTX Packet Transport Routers])

```

user@host> show interfaces detail em0
Physical interface: em0, Enabled, Physical link is Up
 Interface index: 8, SNMP ifIndex: 0, Generation: 3
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking:

```

```

Unspecified,
 Speed: 1000mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 15255909
 Output bytes : 4608
 Input packets : 214753
 Output packets: 72
 IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

 Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 14394630
 Output bytes : 3024
 Input packets : 214723
 Output packets: 72
 Local statistics:
 Input bytes : 14394630
 Output bytes : 3024
 Input packets : 214723
 Output packets: 72
 Protocol inet, MTU: 1500, Generation: 1, Route table: 0
 Flags: Is-Primary
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 192.168.3/24, Local: 192.168.3.30,
 Broadcast: 192.168.3.255, Generation: 1

```

**show interfaces extensive (Management Ethernet)**

```

user@host> show interfaces fxp0 extensive
Physical interface: fxp0, Enabled, Physical link is Up
 Interface index: 1, SNMP ifIndex: 1, Generation: 0
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Half-Duplex
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 6678904
 Output bytes : 169657
 Input packets : 83946
 Output packets: 1127
 Input errors:
 Errors: 12, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
 Policed discards: 0, Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
 Resource errors: 0

Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500, Generation: 6, Route table: 0
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.168.64/21, Local: 192.168.70.143,
 Broadcast: 192.168.71.255, Generation: 1

```

**show interfaces extensive (Management Ethernet [TX Matrix Plus Router])**

```

user@host> show interfaces em0 extensive

```

```

Physical interface: em0, Enabled, Physical link is Up
 Interface index: 8, SNMP ifIndex: 17, Generation: 2
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 127120
 Output bytes : 1357414
 Input packets : 1843
 Output packets: 5372
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed
discards: 0, Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 119748
 Output bytes : 1335719
 Input packets : 1843
 Output packets: 5372
 Local statistics:
 Input bytes : 119748
 Output bytes : 1335719
 Input packets : 1843
 Output packets: 5372
 Protocol inet, MTU: 1500, Generation: 1, Route table: 0

```

```

Flags: Is-Primary
Addresses, Flags: Is-Preferred Is-Primary
 Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast:
192.168.178.127, Generation: 1

```

### show interfaces extensive (Management Ethernet [PTX Series Packet Transport Routers])

```

user@host> show interfaces extensive em0
Physical interface: em0, Enabled, Physical link is Up
 Interface index: 8, SNMP ifIndex: 0, Generation: 3
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking:
Unspecified,
 Speed: 1000mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 15236459
 Output bytes : 4608
 Input packets : 214482
 Output packets: 72
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets: 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
 Policed discards: 0, Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
 Resource errors: 0

 Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
 Flags: SNMP-Traps Encapsulation: ENET2

```

```

Traffic statistics:
 Input bytes : 14376264
 Output bytes : 3024
 Input packets: 214452
 Output packets: 72
Local statistics:
 Input bytes : 14376264
 Output bytes : 3024
 Input packets: 214452
 Output packets: 72
Protocol inet, MTU: 1500, Generation: 1, Route table: 0
 Flags: Is-Primary
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 192.168.3/24, Local: 192.168.3.30,
 Broadcast: 192.168.3.255, Generation: 1

```

### show interfaces mgmtre0 (Management Ethernet [PTX5000 Router])

```

user@host> show interfaces mgmtre0 extensive
Physical interface: mgmtre0, Enabled, Physical link is Up
 Interface index: 1001, SNMP ifIndex: 501
 Link-level type: Ethernet, MTU: 1500
 Device flags : Present
 Interface flags: None
 Link flags : None
 Current address: ec:9e:cd:06:30:da, Hardware address: ec:9e:cd:06:30:da
 Last flapped : Never

Logical interface mgmtre0.0 (Index 1001) (SNMP ifIndex 503)
 Flags: Encapsulation: ENET2
 Protocol inet, MTU: 1486
 Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.92.248/23, Local: 10.92.248.22,
 Broadcast: 10.92.249.255
 Protocol multiservice, MTU: Unlimited
 Flags: None

```

### show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp1 brief
Physical interface: fxp1, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

Logical interface fxp1.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet 10.0.0.4/8
 inet6 fe80::200:ff:fe00:4/64
 fec0::10:0:0:4/64
 tnp 4
```

### show interfaces brief (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

Logical interface em0.0
 Flags: SNMP-Traps Encapsulation: ENET2
 inet 192.168.178.11/25
```

### show interfaces brief (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking:
 Unspecified,
 Speed: 1000mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
```

```

Logical interface em0.0
Flags: SNMP-Traps Encapsulation: ENET2
inet 192.168.3.30/24

root@aboslutely> show interfaces em0 terse
Interface Admin Link Proto Local Remote
em0 up up
em0.0 up up inet 192.168.3.30/24

```

### show interfaces (Internal Ethernet)

```

user@host> show interfaces fxp1
Physical interface: fxp1, Enabled, Physical link is Up
Interface index: 2, SNMP ifIndex: 2
Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
Last flapped : Never
 Input packets : 30655
 Output packets: 33323

Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500
 Flags: Is-Primary
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255
Protocol inet6, MTU: 1500
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred
 Destination: fe80::/64, Local: fe80::200:ff:fe00:4
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: fec0::/64, Local: fec0::10:0:0:4
Protocol tnp, MTU: 1500
 Flags: Primary, Is-Primary
 Addresses
 Local: 4

```

## show interfaces (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0
Physical interface: ixgbe0, Enabled, Physical link is Up
 Interface index: 2, SNMP ifIndex: 116
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
 Last flapped : Never
 Input packets : 2301738
 Output packets: 3951155

Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117)
 Flags: SNMP-Traps Encapsulation: ENET2
 Input packets : 2301595
 Output packets: 3951155
 Protocol inet, MTU: 1500
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred
 Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255
 Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
 Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 192.168.0.4
 Protocol inet6, MTU: 1500
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred
 Destination: fe80::/64, Local: fe80::200:ff:fe22:4
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: fec0::/64, Local: fec0::a:22:0:4
 Protocol tnp, MTU: 1500
 Flags: Primary, Is-Primary
 Addresses
 Local: 0x22000004
```

## show interfaces detail (Internal Ethernet)

```
user@host> show interfaces fxp1 detail
Physical interface: fxp1, Enabled, Physical link is Up
 Interface index: 2, SNMP ifIndex: 2, Generation: 1
```

```

Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 100mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
Alternate link address: Unspecified
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 2339969
 Output bytes : 15880707
 Input packets: 30758
 Output packets: 33443

Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500, Generation: 7, Route table: 1
 Flags: Is-Primary
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
 Generation: 3
 Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred
 Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
 Broadcast: Unspecified, Generation: 5
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
 Generation: 7
 Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
 Flags: Primary, Is-Primary
 Addresses, Flags: None
 Destination: Unspecified, Local: 4, Broadcast: Unspecified,
 Generation: 8

```

**show interfaces detail (Internal Ethernet [TX Matrix Plus Router])**

```
user@host> show interfaces ixgbe0 detail
Physical interface: ixgbe0, Enabled, Physical link is Up
 Interface index: 2, SNMP ifIndex: 116, Generation: 3
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
Alternate link address: Unspecified
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 238172825
 Output bytes : 1338948955
 Input packets: 2360984
 Output packets: 4061512
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
 Flags: SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
 Input bytes : 228720309
 Output bytes : 1261387447
 Input packets: 2360841
 Output packets: 4061512
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Local statistics:
 Input bytes : 228720309
 Output bytes : 1261387447
```

```

Input packets: 2360841
Output packets: 4061512
Protocol inet, MTU: 1500, Generation: 2, Route table: 1
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred
 Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255,
Generation: 2
 Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
 Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255,
Generation: 3
Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred
 Destination: fe80::/64, Local: fe80::200:ff:fe22:4
Generation: 4
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: fec0::/64, Local: fec0::a:22:0:4
Protocol tnp, MTU: 1500, Generation: 5
Generation: 4, Route table: 1
 Flags: Primary, Is-Primary
 Addresses, Flags: None
 Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6

```

### show interfaces extensive (internal Ethernet)

```

user@host> show interfaces fxp1 extensive
Physical interface: fxp1, Enabled, Physical link is Up
 Interface index: 2, SNMP ifIndex: 2, Generation: 1
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
 Speed: 100mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:

```

```

Input bytes : 2349897
Output bytes : 15888605
Input packets: 30896
Output packets: 33607
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
 Policed discards: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
 Resource errors: 0

Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
 Flags: SNMP-Traps Encapsulation: ENET2
 Protocol inet, MTU: 1500, Generation: 7, Route table: 1
 Flags: Is-Primary
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
 Generation: 3
 Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
 Flags: Is-Primary
 Addresses, Flags: Is-Preferred
 Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
 Broadcast: Unspecified, Generation: 5
 Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
 Generation: 7
 Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
 Flags: Primary, Is-Primary
 Addresses, Flags: None
 Destination: Unspecified, Local: 4, Broadcast: Unspecified,
 Generation: 8

```

### show interfaces extensive (internal Ethernet [TX Matrix Plus Router])

```

user@host> show interfaces ixgbe0 extensive
Physical interface: ixgbe0, Enabled, Physical link is Up
 Interface index: 2, SNMP ifIndex: 116, Generation: 3
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
 Device flags : Present Running
 Interface flags: SNMP-Traps

```

```

Link type : Full-Duplex
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
Alternate link address: Unspecified
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 242730780
 Output bytes : 1348312269
 Input packets: 2398737
 Output packets: 4133510
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed
discards: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

```

Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

```

 Input bytes : 233127252
 Output bytes : 1269350897
 Input packets: 2398594
 Output packets: 4133510

```

IPv6 transit statistics:

```

 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

Local statistics:

```

 Input bytes : 233127252
 Output bytes : 1269350897
 Input packets: 2398594
 Output packets: 4133510

```

Protocol inet, MTU: 1500, Generation: 2, Route table: 1

Flags: Is-Primary

```

Addresses, Flags: Is-Preferred
 Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255,
Generation: 2
Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
 Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255,
Generation: 3
Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
 Flags: Is-Primary
Addresses, Flags: Is-Preferred
 Destination: fe80::/64, Local: fe80::200:ff:fe22:4
Generation: 4
Addresses, Flags: Is-Default Is-Preferred Is-Primary
 Destination: fec0::/64, Local: fec0::a:22:0:4
Protocol tnp, MTU: 1500, Generation: 5
Generation: 4, Route table: 1
 Flags: Primary, Is-Primary
Addresses, Flags: None
 Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified,
Generation: 6

```

## Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

## show interfaces (PPPoE)

### IN THIS SECTION

- [Syntax | 1628](#)
- [Description | 1628](#)
- [Options | 1628](#)
- [Required Privilege Level | 1628](#)
- [Output Fields | 1629](#)
- [Sample Output | 1640](#)

## Syntax

```
show interfaces pp0.logical
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Description

(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.

## Options

|                                           |                                                                               |
|-------------------------------------------|-------------------------------------------------------------------------------|
| <b>pp0.<i>logical</i></b>                 | Display standard status information about the PPPoE interface.                |
| <b>brief   detail   extensive   terse</b> | (Optional) Display the specified level of output.                             |
| <b>descriptions</b>                       | (Optional) Display interface description strings.                             |
| <b>media</b>                              | (Optional) Display media-specific information about PPPoE interfaces.         |
| <b>snmp-index <i>snmp-index</i></b>       | (Optional) Display information for the specified SNMP index of the interface. |
| <b>statistics</b>                         | (Optional) Display PPPoE interface statistics.                                |

## Required Privilege Level

view

## Output Fields

Table 158 on page 1629 lists the output fields for the **show interfaces (PPPoE)** command. Output fields are listed in the approximate order in which they appear.

**Table 158: show interfaces (PPPoE) Output Fields**

| Field Name                | Field Description                                                                                                                             | Level of Output                 |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Physical Interface</b> |                                                                                                                                               |                                 |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                               | All levels                      |
| <b>Enabled</b>            | State of the interface. Possible values are described in the “Enabled Field” section under <a href="#">Common Output Fields Description</a> . | All levels                      |
| <b>Interface index</b>    | Physical interface index number, which reflects its initialization sequence.                                                                  | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                 | <b>detail extensive</b><br>none |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                             | <b>detail extensive</b>         |
| <b>Type</b>               | Physical interface type (PPPoE).                                                                                                              | All levels                      |
| <b>Link-level type</b>    | Encapsulation on the physical interface (PPPoE).                                                                                              | All levels                      |
| <b>MTU</b>                | MTU size on the physical interface.                                                                                                           | All levels                      |
| <b>Clocking</b>           | Reference clock source. It can be <b>Internal</b> or <b>External</b> .                                                                        | All levels                      |
| <b>Speed</b>              | Speed at which the interface is running.                                                                                                      | All levels                      |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name                    | Field Description                                                                                                                                           | Level of Output         |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Device flags</b>           | Information about the physical device. Possible values are described in the “Device Flags” section under <a href="#">Common Output Fields Description</a> . | All levels              |
| <b>Interface flags</b>        | Information about the interface. Possible values are described in the “Interface Flags” section under <a href="#">Common Output Fields Description</a> .    | All levels              |
| <b>Link type</b>              | Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .                                                                                    | All levels              |
| <b>Link flags</b>             | Information about the interface. Possible values are described in the “Link Flags” section under <a href="#">Common Output Fields Description</a> .         | All levels              |
| <b>Input rate</b>             | Input rate in bits per second (bps) and packets per second (pps).                                                                                           | None specified          |
| <b>Output rate</b>            | Output rate in bps and pps.                                                                                                                                 | None specified          |
| <b>Physical Info</b>          | Physical interface information.                                                                                                                             | All levels              |
| <b>Hold-times</b>             | Current interface hold-time up and hold-time down, in milliseconds.                                                                                         | <b>detail extensive</b> |
| <b>Current address</b>        | Configured MAC address.                                                                                                                                     | <b>detail extensive</b> |
| <b>Hardware address</b>       | MAC address of the hardware.                                                                                                                                | <b>detail extensive</b> |
| <b>Alternate link address</b> | Backup address of the link.                                                                                                                                 | <b>detail extensive</b> |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output         |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                        | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <p><b>NOTE:</b> These fields include dropped traffic and exception traffic, as those fields are not separately defined.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b> |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output  |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Input errors</b> | <p>Input errors on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of incoming frame terminations and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of B chip Tx drops and IXP Tx net transmit drops.</li> </ul> | <b>extensive</b> |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output                 |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Output errors</b>     | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b> –Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), then the cable, the far-end system, or the PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminations and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of B chip Tx drops and IXP Tx net transmit drops.</li> </ul> | <b>extensive</b>                |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                 |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels                      |
| <b>Index</b>             | Logical interface index number (which reflects its initialization sequence).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>      | Logical interface SNMP interface index number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b><br>none |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Generation</b>     | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <b>Flags</b>          | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <a href="#">Common Output Fields Description</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | All levels              |
| <b>Encapsulation</b>  | Type of encapsulation configured on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels              |
| <b>PPP parameters</b> | <p>PPP status:</p> <ul style="list-style-type: none"> <li>• LCP restart timer—Length of time (in milliseconds) between successive Link Control Protocol (LCP) configuration requests.</li> <li>• NCP restart timer—Length of time (in milliseconds) between successive Network Control Protocol (NCP) configuration requests.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>           |
| <b>PPPoE</b>          | <p>PPPoE status:</p> <ul style="list-style-type: none"> <li>• <b>State</b>—State of the logical interface (<b>up</b> or <b>down</b>).</li> <li>• <b>Session ID</b>—PPPoE session ID.</li> <li>• <b>Service name</b>—Type of service required. Can be used to indicate an Internet service provider (ISP) name or a class or quality of service.</li> <li>• <b>Configured AC name</b>—Configured access concentrator name.</li> <li>• <b>Auto-reconnect timeout</b>—Time after which to try to reconnect after a PPPoE session is terminated, in seconds.</li> <li>• <b>Idle Timeout</b>—Length of time (in seconds) that a connection can be idle before disconnecting.</li> <li>• <b>Underlying interface</b>—Interface on which PPPoE is running.</li> </ul> | All levels              |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output         |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Link</b>                    | Name of the physical interfaces for member links in an aggregated Ethernet bundle for a PPPoE over aggregated Ethernet configuration. PPPoE traffic goes out on these interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                    | All levels              |
| <b>Traffic statistics</b>      | Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.                                                                                                                                                                                                                                                                                           | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b> |
| <b>Local statistics</b>        | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b> |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output         |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Transit statistics</b> | <p>Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. This counter usually takes less than 1 second to stabilize.</p> <p><b>NOTE:</b> The packet and byte counts in these fields include traffic that is dropped and does not leave the router.</p>                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>Keepalive settings</b> | <p>(PPP and HDLC) Configured settings for keepalives.</p> <ul style="list-style-type: none"> <li>• <b>interval <i>seconds</i></b>—The time in seconds between successive keepalive requests. The range is <b>10</b> seconds through <b>32,767</b> seconds, with a default of <b>10</b> seconds.</li> <li>• <b>down-count <i>number</i></b>—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is <b>1</b> through <b>255</b>, with a default of <b>3</b>.</li> <li>• <b>up-count <i>number</i></b>—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is <b>1</b> through <b>255</b>, with a default of <b>1</b>.</li> </ul> | <b>detail extensive</b> |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output              |
|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Keepalive statistics</b> | <p>(PPP and HDLC) Information about keepalive packets.</p> <ul style="list-style-type: none"> <li>• <b>Input</b>—Number of keepalive packets received by PPP. <ul style="list-style-type: none"> <li>• <b>(last seen 00:00:00 ago)</b>—Time the last keepalive packet was received, in the format <i>hh:mm:ss</i>.</li> </ul> </li> <li>• <b>Output</b>—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received. <ul style="list-style-type: none"> <li>• <b>(last seen 00:00:00 ago)</b>—Time the last keepalive packet was sent, in the format <i>hh:mm:ss</i>.</li> </ul> </li> </ul> <p>(MX Series routers with MPCs/MICs) When an MX Series router with MPCs/MICs is using PPP fast keepalive for a PPP link, the display does not include the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.</p> | <b>detail extensive</b>      |
| <b>Input packets</b>        | Number of packets received on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | None specified               |
| <b>Output packets</b>       | Number of packets transmitted on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | None specified               |
| <b>LCP state</b>            | <p>(PPP) Link Control Protocol state.</p> <ul style="list-style-type: none"> <li>• <b>Conf-ack-received</b>—Acknowledgement was received.</li> <li>• <b>Conf-ack-sent</b>—Acknowledgement was sent.</li> <li>• <b>Conf-req-sent</b>—Request was sent.</li> <li>• <b>Down</b>—LCP negotiation is incomplete (not yet completed or has failed).</li> <li>• <b>Not-configured</b>—LCP is not configured on the interface.</li> <li>• <b>Opened</b>—LCP negotiation is successful.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>none detail extensive</b> |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output                 |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>NCP state</b>  | <p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> <li>• <b>Conf-ack-received</b>—Acknowledgement was received.</li> <li>• <b>Conf-ack-sent</b>—Acknowledgement was sent.</li> <li>• <b>Conf-req-sent</b>—Request was sent.</li> <li>• <b>Down</b>—NCP negotiation is incomplete (not yet completed or has failed).</li> <li>• <b>Not-configured</b>—NCP is not configured on the interface.</li> <li>• <b>Opened</b>—NCP negotiation is successful.</li> </ul>                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b><br>none |
| <b>CHAP state</b> | <p>(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction.</p> <ul style="list-style-type: none"> <li>• <b>Chap-Chal-received</b>—Challenge was received but response not yet sent.</li> <li>• <b>Chap-Chal-sent</b>—Challenge was sent.</li> <li>• <b>Chap-Resp-received</b>—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.)</li> <li>• <b>Chap-Resp-sent</b>—Response was sent for the challenge received.</li> <li>• <b>Closed</b>—CHAP authentication is incomplete.</li> <li>• <b>Failure</b>—CHAP authentication failed.</li> <li>• <b>Not-configured</b>—CHAP is not configured on the interface.</li> <li>• <b>Success</b>—CHAP authentication was successful.</li> </ul> | none <b>detail extensive</b>    |
| <b>Protocol</b>   | Protocol family configured on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b><br>none |

Table 158: show interfaces (PPPoE) Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                           | Level of Output                 |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <i>protocol-family</i>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                    | <b>brief</b>                    |
| <b>MTU</b>              | MTU size on the logical interface.                                                                                                                                                          | <b>detail extensive</b><br>none |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.                                                                                                                           | <b>detail extensive</b>         |
| <b>Route table</b>      | Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .                                                          | <b>detail extensive</b><br>none |
| <b>Flags</b>            | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <a href="#">Common Output Fields Description</a> .                           | <b>detail extensive</b><br>none |
| <b>Addresses, Flags</b> | Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under <a href="#">Common Output Fields Description</a> . | <b>detail extensive</b><br>none |
| <b>Destination</b>      | IP address of the remote side of the connection.                                                                                                                                            | <b>detail extensive</b><br>none |
| <b>Local</b>            | IP address of the logical interface.                                                                                                                                                        | <b>detail extensive</b><br>none |
| <b>Broadcast</b>        | Broadcast address.                                                                                                                                                                          | <b>detail extensive</b><br>none |

## Sample Output

### show interfaces (PPPoE)

```

user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 24
 Type: PPPoE, Link-level type: PPPoE, MTU: 1532
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps
 Link type : Full-Duplex
 Link flags : None
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)

Logical interface pp0.0 (Index 72) (SNMP ifIndex 72)
 Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
 PPPoE:
 State: SessionDown, Session ID: None,
 Service name: None, Configured AC name: sapphire,
 Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
 Underlying interface: at-5/0/0.0 (Index 70)
 Input packets : 0
 Output packets: 0
 LCP state: Not-configured
 NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
 mpls: Not-configured
 CHAP state: Closed
 Protocol inet, MTU: 100
 Flags: User-MTU, Negotiate-Address

```

### show interfaces (PPPoE over Aggregated Ethernet)

```

user@host> show interfaces pp0.1073773821
Logical interface pp0.1073773821 (Index 80) (SNMP ifIndex 32584)
 Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
 PPPoE:
 State: SessionUp, Session ID: 1,
 Session AC name: alcor, Remote MAC address: 00:00:5e:00:53:01,
 Underlying interface: demux0.100 (Index 88)

```

```

Link:
 ge-1/0/0.32767
 ge-1/0/1.32767
Input packets : 6
Output packets: 6
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Success
 Protocol inet, MTU: 1500
 Flags: Sendbcast-pkt-to-re
 Addresses, Flags: Is-Primary
 Local: 203.0.113.1

```

### show interfaces brief (PPPoE)

```

user@host> show interfaces pp0 brief
Physical interface: pp0, Enabled, Physical link is Up
 Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps

Logical interface pp0.0
 Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
 PPPoE:
 State: SessionDown, Session ID: None,
 Service name: None, Configured AC name: sapphire,
 Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
 Underlying interface: at-5/0/0.0 (Index 70)
 inet

```

### show interfaces detail (PPPoE)

```

user@host> show interfaces pp0 detail
Physical interface: pp0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 24, Generation: 9
 Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
 Device flags : Present Running

```

```

Interface flags: Point-To-Point SNMP-Traps
Link type : Full-Duplex
Link flags : None
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
Logical interface pp0.0 (Index 72) (SNMP ifIndex 72) (Generation 14)
 Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
 PPPoE:
 State: SessionDown, Session ID: None,
 Service name: None, Configured AC name: sapphire,
 Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
 Underlying interface: at-5/0/0.0 (Index 70)
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Local statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
 Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
LCP state: Not-configured
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Closed
 Protocol inet, MTU: 100, Generation: 14, Route table: 0
 Flags: User-MTU, Negotiate-Address

```

**show interfaces extensive (PPPoE on M120 and M320 Routers)**

```

user@host> show interfaces pp0 extensive
Physical interface: pp0, Enabled, Physical link is Up
 Interface index: 128, SNMP ifIndex: 93, Generation: 129
 Type: PPPoE, Link-level type: PPPoE, MTU: 1532, Speed: Unspecified
 Device flags : Present Running
 Interface flags: Point-To-Point SNMP-Traps
 Link type : Full-Duplex
 Link flags : None
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: Unspecified, Hardware address: Unspecified
 Alternate link address: Unspecified
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 972192 0 bps
 Output bytes : 975010 0 bps
 Input packets : 1338 0 pps
 Output packets : 1473 0 pps
 IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets : 0
 Output packets : 0
 Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed
discards: 0,
 Resource errors: 0
 Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface pp0.0 (Index 69) (SNMP ifIndex 96) (Generation 194)
 Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPPoE
 PPPoE:
 State: SessionUp, Session ID: 26,
 Session AC name: None, AC MAC address: 00:00:5e:00:53:12,
 Service name: None, Configured AC name: None,
 Auto-reconnect timeout: Never, Idle timeout: Never,
 Underlying interface: ge-3/0/1.0 (Index 67)
 Traffic statistics:

```

```

Input bytes : 252
Output bytes : 296
Input packets: 7
Output packets: 8
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Local statistics:
 Input bytes : 252
 Output bytes : 296
 Input packets: 7
 Output packets: 8
Transit statistics:
 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
 Input : 1 (last seen 00:00:00 ago)
 Output: 1 (last sent 00:00:03 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-
configured
CHAP state: Closed
PAP state: Closed
Protocol inet, MTU: 1492, Generation: 171, Route table: 0
 Flags: None
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified,
Generation: 206

```

## Release Information

Command introduced before Junos OS Release 7.4.

## show interfaces interface-set (Ethernet Interface Set)

### IN THIS SECTION

- [Syntax | 1645](#)
- [Description | 1645](#)
- [Options | 1645](#)
- [Required Privilege Level | 1645](#)
- [Output Fields | 1646](#)
- [Sample Output | 1648](#)
- [Release Information | 1651](#)

### Syntax

```
show interfaces interface-set interface-set-name
<detail | terse>
```

### Description

Display information about the specified gigabit or 10-Gigabit Ethernet interface set.

You can also use the **show interfaces interface-set** command to display information about agent circuit identifier (ACI) interface sets.

### Options

**interface-set *interface-set-name***      Display information about the specified Gigabit Ethernet, 10-Gigabit Ethernet, ACI, or ALI interface set.

**detail | terse**      (Optional) Display the specified level of output.

### Required Privilege Level

view

## Output Fields

Table 159 on page 1646 describes the information for the **show interfaces interface-set** command. Output fields are listed in the approximate order in which they appear.

**Table 159: Ethernet show interfaces interface-set Output Fields**

| Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Level of Output    |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>Physical Interface</b>  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                    |
| <b>Interface set</b>       | <p>Name of the interface set or sets.</p> <p>For ACI interface sets, the set name is prefixed with <b>aci-</b>.</p> <p>For ALI interface sets, the set name is prefixed with the trusted option that the interface set is based on:</p> <ul style="list-style-type: none"> <li>• <b>aci-</b>— The ACI is configured as the trusted option.</li> <li>• <b>ari-</b>— The ARI is configured as the trusted option.</li> <li>• <b>aci+ari-</b>— Both ACI and ARI are configured as the trusted option.</li> <li>• <b>noids-</b>— Neither the ACI nor the ARI is configured as the trusted option and neither ACI nor ARI is received.</li> </ul> | All levels         |
| <b>Interface set index</b> | Index number of the interface set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail none</b> |
| <b>ACI VLAN</b>            | <p>For ACI interface sets, the string received in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and the DSL line on the access node. Only the <b>Agent Circuit ID</b> can be used to create the interface set.</p> <p><b>NOTE:</b> The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the <b>line-identity</b> autoconfiguration stanza.</p>                                                                                                                                                                                                     | <b>detail none</b> |

Table 159: Ethernet show interfaces interface-set Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>Line Identity</b>           | <p>For ALI interface sets, the trusted option received in DHCP or PPPoE control packets that uniquely identifies the subscriber's access node and the DSL line on the access node. The trusted option can be either or both of the following:</p> <ul style="list-style-type: none"> <li>• <b>Agent Circuit ID</b>—The ACI value</li> <li>• <b>Agent Remote ID</b>—The ARI value.</li> </ul> <p><b>NOTE:</b> When only <b>accept-no-ids</b> is configured as the trusted option, this field is not displayed.</p> <p><b>NOTE:</b> The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the <b>agent-circuit-id</b> autoconfiguration stanza.</p> | <b>detail none</b> |
| <b>PPPoE</b>                   | Dynamic PPPoE subscriber interface that the router creates using the ACI or ALI interface set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail none</b> |
| <b>Max Sessions</b>            | For dynamic PPPoE subscriber interfaces, maximum number of PPPoE logical interfaces that that can be activated on the underlying interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail none</b> |
| <b>Max Sessions VSA Ignore</b> | For dynamic PPPoE subscriber interfaces, whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configured with the <b>max-sessions</b> statement: <b>Off</b> (default) or <b>On</b> .                                                                                                                                                                                                                                                                                                            | <b>detail none</b> |

Table 159: Ethernet show interfaces interface-set Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes and number of bytes per second received and transmitted on the interface set</li> <li>• <b>Input packets, Output packets</b>—Number of packets and number of packets per second received and transmitted on the interface set.</li> </ul> | <b>detail</b>      |
| <b>Egress queues supported</b> | Total number of egress queues supported on the specified interface set.                                                                                                                                                                                                                                                                                                                                                                  | <b>detail</b>      |
| <b>Egress queues in use</b>    | Total number of egress queues used on the specified interface set.                                                                                                                                                                                                                                                                                                                                                                       | <b>detail</b>      |
| <b>Queue counters</b>          | <b>Queued packets, Transmitted packets, and Dropped packets</b> statistics for the four forwarding classes.                                                                                                                                                                                                                                                                                                                              | <b>detail</b>      |
| <b>Members</b>                 | List of all interface sets or, for ACI interface sets, list of all subscriber interfaces belonging to the specified ACI interface set.                                                                                                                                                                                                                                                                                                   | <b>detail none</b> |

## Sample Output

### show interfaces interface-set terse

```

user@host> show interfaces interface-set terse
Interface set:
 iflset-xe-11/3/0-0
 ge-1/0/1-0
 ge-1/0/1-2

```

**show interfaces interface-set detail**

```

user@host> show interfaces interface-set iflset-xe-11/3/0-0 detail
Interface set: iflset-xe-11/3/0-0
Interface set index: 19
Traffic statistics:
 Output bytes : 751017840 401673504 bps
 Output packets: 11044380 738377 pps
Egress queues: 4 supported, 4 in use
Queue counters:
 Queued packets Transmitted packets Dropped packets
 0 211091327 11044380 199995746
 1 0 0 0
 2 0 0 0
 3 0 0 0
Members:
 xe-11/3/0.0

```

**show interfaces interface-set (ACI Interface Set based on ACI)**

```

user@host> show interfaces interface-set
Interface set: aci-1001-ge-5/2/0.10
Interface set index: 1
Interface set snmp index: 67108865
ACI VLAN:
 Agent Circuit ID: circuit0
PPPoE:
 Max Sessions: 32000, Max Sessions VSA Ignore: Off
Members:
 demux0.3221225472

```

**show interfaces interface-set (ACI Interface Set based on ACI Trusted Option)**

```

user@host> show interfaces interface-set
Interface set: ari-1002-demux0.3221225473
Interface set index: 2
Interface set snmp index: 67108866
Line Identity:
 Agent Circuit ID: remote20
PPPoE:

```

```

 Max Sessions: 32000, Max Sessions VSA Ignore: Off
Members:
 demux0.3221225474

```

### show interfaces interface-set (ACI Interface Set based on ARI Trusted Option)

```

user@host> show interfaces interface-set
Interface set: aci-1002-demux0.3221225473
Interface set index: 2
Interface set snmp index: 67108866
Line Identity:
 Agent Remote ID: remote20
PPPoE:
 Max Sessions: 32000, Max Sessions VSA Ignore: Off
Members:
 demux0.3221225474

```

### show interfaces interface-set (ACI Interface Set based on ARI Trusted Option when both ACI and ARI are received)

```

user@host> show interfaces interface-set
Interface set: ari-1002-demux0.3221225473
Interface set index: 2
Interface set snmp index: 67108866
Line Identity:
 Agent Remote ID: remote20
PPPoE:
 Max Sessions: 32000, Max Sessions VSA Ignore: Off
Members:
 demux0.3221225474

```

### show interfaces interface-set (ACI Interface Set based on Accept-No-IDs Trusted Option when neither ACI nor ARI is received)

```

user@host> show interfaces interface-set
Interface set: noids-1002-demux0.3221225473
Interface set index: 2
Interface set snmp index: 67108866

```

```
Members:
 demux0.3221225474
```

## show interfaces interface-set (L2BSA and PPPoE Subscribers)

```
user@host> show interfaces interface-set
Interface set: ge-1/0/4
 Interface set index: 6
 Members:
 ge-1/0/4.1073741908
 pp0.1073741907
```

## Release Information

Command introduced in Junos OS Release 8.5.

## RELATED DOCUMENTATION

*Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

*Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

## show interfaces interface-set queue

### IN THIS SECTION

- [Syntax | 1652](#)
- [Description | 1652](#)
- [Options | 1652](#)
- [Required Privilege Level | 1652](#)
- [Output Fields | 1652](#)
- [Sample Output | 1654](#)
- [Release Information | 1662](#)

## Syntax

```
show interfaces interface-set queue interface-set-name
<aggregate | remaining-traffic>
<forwarding-class class-name>
```

## Description

Display information about the gigabit or 10-Gigabit Ethernet interface set queue. Supported in MX Series routers with enhanced queuing DPCs.

## Options

|                                           |                                                                                                                                                         |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b><i>interface-set-name</i></b>          | (Optional) Display information about the specified gigabit or 10-Gigabit Ethernet interface set. Wildcard values can be used in the interface set name. |
| <b>aggregate</b>                          | (Optional) Display the aggregated queuing statistics of all member logical interfaces for interface sets that have traffic-control profiles configured. |
| <b>both-ingress-egress</b>                | (Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics.                                      |
| <b>egress</b>                             | (Optional) Display egress queue statistics.                                                                                                             |
| <b>forwarding-class <i>class-name</i></b> | (Optional) Display queuing statistics for the specified forwarding class.                                                                               |
| <b>ingress</b>                            | (Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics.                                                                              |
| <b>remaining-traffic</b>                  | (Optional) Display the queuing statistics of all member logical interfaces for interface sets that do not have traffic-control profiles configured.     |

## Required Privilege Level

view

## Output Fields

[Table 160 on page 1653](#) describes the information for the **show interfaces interface-set queue** command.

Table 160: Ethernet show interfaces interface-set queue Output Fields

| Field Name                          | Field Description                                                            | Level of Output |
|-------------------------------------|------------------------------------------------------------------------------|-----------------|
| <b>Physical Interface</b>           |                                                                              |                 |
| <b>Interface set</b>                | Name of the interface set.                                                   | All levels      |
| <b>Interface set index</b>          | Index number of the interface set.                                           | All levels      |
| <b>Forwarding classes supported</b> | Total number of forwarding classes supported on the specified interface set. | All levels      |
| <b>Forwarding classes in use</b>    | Total number of forwarding classes used on the specified interface set.      | All levels      |
| <b>Egress queues supported</b>      | Total number of egress queues supported on the specified interface set.      | All levels      |
| <b>Egress queues in use</b>         | Total number of egress queues used on the specified interface set.           | All levels      |
| <b>Ingress queues supported</b>     | Total number of ingress queues supported on the specified interface set.     | All levels      |
| <b>Ingress queues in use</b>        | Total number of ingress queues used on the specified interface set.          | All levels      |
| <b>Queue</b>                        | Egress or ingress queue number for the statistics being displayed.           | All levels      |
| <b>Forwarding classes</b>           | Forwarding class name for the statistics being displayed.                    | All levels      |

Table 160: Ethernet show interfaces interface-set queue Output Fields (Continued)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Queued</b>      | <p><b>Packet</b> and <b>Byte</b> statistics for the specified queue.</p> <ul style="list-style-type: none"> <li>• <b>Packets</b>—Number of packets queued and input rate in packets per second.</li> <li>• <b>Bytes</b>—Number of bytes queued and input rate in bytes per second.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| <b>Transmitted</b> | <p><b>Packet</b> and <b>Byte</b> statistics for the specified forwarding class.</p> <ul style="list-style-type: none"> <li>• <b>Packets</b>—Number of packets transmitted and transmit rate in packets per second.</li> <li>• <b>Bytes</b>—Number of bytes transmitted and transmit rate in bytes per second.</li> <li>• <b>Tail-dropped packets</b>—Number of packets tail dropped.</li> <li>• <b>RED-dropped packets</b>—Number of RED-dropped packets for the <b>low</b>, <b>medium-low</b>, <b>medium-high</b>, and <b>high</b> loss priorities.</li> <li>• <b>RED-dropped bytes</b>—Number of RED-dropped bytes for the <b>low</b>, <b>medium-low</b>, <b>medium-high</b>, and <b>high</b> loss priorities.</li> </ul> | All levels      |

## Sample Output

### show interfaces interface-set queue (Gigabit Ethernet)

```

user@host> show interfaces interface-set queue ge-2/2/0-0
Interface set: ge-2/2/0-0
 Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 3998482 1 pps
 Bytes : 271896884 688 bps

```

```

Transmitted:
 Packets : 1077474 1 pps
 Bytes : 73268340 688 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 2921008 0 pps
 Low : 2921008 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 198628544 0 bps
 Low : 198628544 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 ...

```

### show interfaces interface-set queue both-ingress-egress (Enhanced DPC)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 both-ingress-egress
Interface set: ge-2/2/0-0
 Interface set index: 3
 Forwarding classes: 16 supported, 4 in use
 Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 185968478 473161 pps
 Bytes : 10042313520 204441336 bps
 Transmitted:
 Packets : 5441673 13780 pps
 Bytes : 293850342 5952960 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 180526772 459372 pps
 RED-dropped bytes : 9748446282 198451512 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps

```

```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 522021472 473602 pps
Bytes : 28190332480 204599944 bps
Transmitted:
Packets : 5791772 4055 pps
Bytes : 312755688 1751976 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 516227139 469546 pps
RED-dropped bytes : 27876265560 202843872 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 5417304 13797 pps
Bytes : 368429508 7506096 bps
Transmitted:
Packets : 5014996 12769 pps
Bytes : 341019728 6946560 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 402189 1028 pps
 Low : 402189 1028 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 27348852 559536 bps

```

|             |   |          |            |
|-------------|---|----------|------------|
| Low         | : | 27348852 | 559536 bps |
| Medium-low  | : | 0        | 0 bps      |
| Medium-high | : | 0        | 0 bps      |
| High        | : | 0        | 0 bps      |

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 2, Forwarding classes: assured-forwarding

Queued:

|         |   |           |             |
|---------|---|-----------|-------------|
| Packets | : | 5770534   | 3963 pps    |
| Bytes   | : | 396943252 | 2156144 bps |

Transmitted:

|                      |   |           |             |
|----------------------|---|-----------|-------------|
| Packets              | : | 3945152   | 1457 pps    |
| Bytes                | : | 268270336 | 792608 bps  |
| Tail-dropped packets | : | 0         | 0 pps       |
| RED-dropped packets  | : | 1815141   | 2506 pps    |
| Low                  | : | 1815141   | 2506 pps    |
| Medium-low           | : | 0         | 0 pps       |
| Medium-high          | : | 0         | 0 pps       |
| High                 | : | 0         | 0 pps       |
| RED-dropped bytes    | : | 123429524 | 1363536 bps |
| Low                  | : | 123429524 | 1363536 bps |
| Medium-low           | : | 0         | 0 bps       |
| Medium-high          | : | 0         | 0 bps       |
| High                 | : | 0         | 0 bps       |

Queue: 3, Forwarding classes: network-control

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

### show interfaces interface-set queue egress (Enhanced DPC)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 egress
Interface set: ge-2/2/0-0
 Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 3958253 13822 pps
 Bytes : 269217592 7519712 bps
 Transmitted:
 Packets : 3665035 12729 pps
 Bytes : 249222380 6924848 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 293091 1093 pps
 Low : 293091 1093 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 19930188 594864 bps
 Low : 19930188 594864 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 2, Forwarding classes: assured-forwarding

Queued:

|         |   |           |             |
|---------|---|-----------|-------------|
| Packets | : | 5350989   | 3904 pps    |
| Bytes   | : | 368412924 | 2124048 bps |

Transmitted:

|                      |   |           |             |
|----------------------|---|-----------|-------------|
| Packets              | : | 3790469   | 1465 pps    |
| Bytes                | : | 257751892 | 796960 bps  |
| Tail-dropped packets | : | 0         | 0 pps       |
| RED-dropped packets  | : | 1550282   | 2439 pps    |
| Low                  | : | 1550282   | 2439 pps    |
| Medium-low           | : | 0         | 0 pps       |
| Medium-high          | : | 0         | 0 pps       |
| High                 | : | 0         | 0 pps       |
| RED-dropped bytes    | : | 105419176 | 1327088 bps |
| Low                  | : | 105419176 | 1327088 bps |
| Medium-low           | : | 0         | 0 bps       |
| Medium-high          | : | 0         | 0 bps       |
| High                 | : | 0         | 0 bps       |

Queue: 3, Forwarding classes: network-control

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

```

Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
 Low : 0 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
RED-dropped bytes : 0 0 bps
 Low : 0 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

### show interfaces interface-set queue forwarding-class (Gigabit Ethernet)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 forwarding-class best-effort
Interface set: ge-2/2/0-0
 Interface set index: 3
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 101857694 1420083 pps
 Bytes : 6927234456 772532320 bps
 Transmitted:
 Packets : 3984693 55500 pps
 Bytes : 270959592 30192512 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 97870952 1364583 pps
 Low : 97870952 1364583 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 6655225776 742339808 bps
 Low : 6655225776 742339808 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

**show interfaces interface-set queue (Enhanced DPC)**

```

user@host> show interfaces interface-set queue ge-2/2/0-0 ingress
Interface set: foo
 Interface set index: 3
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 149036817 473711 pps
 Bytes : 8048003934 204642936 bps
 Transmitted:
 Packets : 4360749 13891 pps
 Bytes : 235480446 6000912 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 144676035 459820 pps
 RED-dropped bytes : 7812506592 198642024 bps
Queue: 1, Forwarding classes: expedited-forwarding
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 0 0 pps
 RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
 Queued:
 Packets : 485089207 473605 pps
 Bytes : 26195987476 204597576 bps
 Transmitted:
 Packets : 5480799 3959 pps
 Bytes : 295963146 1710504 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 479605853 469646 pps
 RED-dropped bytes : 25898716170 202887072 bps
Queue: 3, Forwarding classes: network-control
 Queued:
 Packets : 0 0 pps
 Bytes : 0 0 bps
 Transmitted:

```

```

Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

### show interfaces interface-set queue remaining-traffic (Gigabit Ethernet)

```

user@host> show interfaces interface-set queue ge-2/2/0-0 remaining-traffic
Interface set: ge-2/2/0-0
 Interface set index: 12
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
 Queued:
 Packets : 2201552 0 pps
 Bytes : 149705536 0 bps
 Transmitted:
 Packets : 609765 0 pps
 Bytes : 41464020 0 bps
 Tail-dropped packets : 0 0 pps
 RED-dropped packets : 1591787 0 pps
 Low : 1591787 0 pps
 Medium-low : 0 0 pps
 Medium-high : 0 0 pps
 High : 0 0 pps
 RED-dropped bytes : 108241516 0 bps
 Low : 108241516 0 bps
 Medium-low : 0 0 bps
 Medium-high : 0 0 bps
 High : 0 0 bps

```

### Release Information

Command introduced in Junos OS Release 8.5.

## show interfaces interval

### IN THIS SECTION

- [Syntax | 1663](#)
- [Description | 1663](#)
- [Options | 1663](#)
- [Required Privilege Level | 1663](#)
- [Output Fields | 1664](#)
- [Sample Output | 1665](#)
- [Release Information | 1667](#)

### Syntax

```
show interfaces interval
<interface-name>
```

### Description

Display the channel service unit (CSU) interface alarm and error count in 15-minute intervals for the past 24 hours. If the system has been operational for less than 24 hours, the maximum number of intervals available is displayed.

### Options

*interface-name* (Optional) Name of a particular interface.

### Required Privilege Level

view

## Output Fields

Table 161 on page 1664 lists the output fields for the **show interfaces interval** command. Output fields are listed in the approximate order in which they appear.

**Table 161: show interfaces interval Output Fields**

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Physical interface</b>              | Name of the interface.                                                                                                                                                                                                                                                                                                                                               |
| <b>SNMP ifIndex</b>                    | SNMP index number for the physical interface.                                                                                                                                                                                                                                                                                                                        |
| <b><i>hh:mm-current</i></b>            | Time of day (in hours and minutes) at the beginning of the latest counter interval. The value of the latest counter interval is always less than 15 minutes.                                                                                                                                                                                                         |
| <b><i>hh:mm-hh:mm</i></b>              | Time of day (in hours and minutes) at the beginning and end of each 15-minute interval.                                                                                                                                                                                                                                                                              |
| <b><i>alarm</i> or <i>event: n</i></b> | Count of alarms and events within each 15-minute interval. The specific alarm or event depends on the interface media type. For a description of the alarm or event listed, see the <b><i>interface-type media</i></b> field (for example, <b>T1 media</b> ) under the <b>show interfaces</b> command for the particular interface type in which you are interested. |
| <b>Interval Total</b>                  | Sum of all the alarm and defect counters for the last 24-hour period.                                                                                                                                                                                                                                                                                                |
| <b>Interval Total</b>                  | Sum of all the alarm and defect counters for the last 24-hour period.                                                                                                                                                                                                                                                                                                |
| <b>Current Day Interval Total</b>      | Sum of all the alarm and defect counters in the current day.<br><b>NOTE:</b> The <b>Current Day Interval</b> output field is reset after 24 hours.                                                                                                                                                                                                                   |
| <b>Previous Day Interval Total</b>     | Sum of all the alarm and defect counters in the previous day.                                                                                                                                                                                                                                                                                                        |

## Sample Output

### show interfaces interval (Channelized OC12)

```

user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
 17:43-current:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 17:28-17:43:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 17:13-17:28:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 16:58-17:13:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 16:43-16:58:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 ...
Interval Total:
 LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
 CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

### show interfaces interval (E3)

```

user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
 17:43-current:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 17:28-17:43:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 17:13-17:28:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
 SEFS: 0, UAS: 0
 16:58-17:13:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,

```

```

SEFS: 0, UAS: 0
16:43-16:58:
 LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
 LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
 CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

## show interfaces interval (SONET/SDH)

```

user@host> show interfaces interval so-2/2/0
Physical interface: so-2/2/0, SNMP ifIndex: 553
 02:53-current:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 02:38-02:53:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 02:23-02:38:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 02:08-02:23:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 01:53-02:08:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 01:38-01:53:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 01:23-01:38:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 01:08-01:23:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 00:53-01:08:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
 00:38-00:53:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:

```

```
0, UAS-P: 0
....
Current Day Interval Total:
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
Previous Day Interval Total (Last updated at 02:23):
 ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P:
0, UAS-P: 0
```

## Release Information

Command introduced before Junos OS Release 7.4.

## RELATED DOCUMENTATION

| [clear interfaces interval](#) | [1231](#)

## show interfaces irb

### IN THIS SECTION

- [Syntax](#) | [1668](#)
- [Description](#) | [1668](#)
- [Options](#) | [1668](#)
- [Additional Information](#) | [1668](#)
- [Required Privilege Level](#) | [1668](#)
- [Output Fields](#) | [1669](#)
- [Sample Output](#) | [1676](#)
- [Release Information](#) | [1680](#)

## Syntax

```
show interfaces irb
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

## Description

Display integrated routing and bridging interfaces information.

## Options

|                                           |                                                                                 |
|-------------------------------------------|---------------------------------------------------------------------------------|
| <b>brief   detail   extensive   terse</b> | (Optional) Display the specified level of output.                               |
| <b>descriptions</b>                       | (Optional) Display interface description strings.                               |
| <b>mac</b>                                | Display hardware MAC address                                                    |
| <b>media</b>                              | (Optional) Display media-specific information about network interfaces.         |
| <b>snmp-index <i>snmp-index</i></b>       | (Optional) Display information for the interface with the specified SNMP index. |
| <b>statistics</b>                         | (Optional) Display static interface statistics.                                 |

## Additional Information

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

## Required Privilege Level

view

## Output Fields

Table 162 on page 1669 lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

**Table 162: show interfaces irb Output Fields**

| Field Name                | Field Description                                                                                                                             | Level of Output                              |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <b>Physical Interface</b> |                                                                                                                                               |                                              |
| <b>Physical interface</b> | Name of the physical interface.                                                                                                               | All levels                                   |
| <b>Enabled</b>            | State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels                                   |
| <b>Proto</b>              | Protocol configured on the interface.                                                                                                         | <b>terse</b>                                 |
| <b>Interface index</b>    | Physical interface index number, which reflects its initialization sequence.                                                                  | <b>detail extensive</b><br>none              |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                 | <b>detail extensive</b><br>none              |
| <b>Type</b>               | Physical interface type.                                                                                                                      | <b>detail extensive</b><br>none              |
| <b>Link-level type</b>    | Encapsulation being used on the physical interface.                                                                                           | <b>detail extensive</b><br><b>brief</b> none |
| <b>MTU</b>                | MTU size on the physical interface.                                                                                                           | <b>detail extensive</b><br><b>brief</b> none |

Table 162: show interfaces irb Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                  | Level of Output                        |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| <b>Clocking</b>         | Reference clock source: <b>Internal</b> or <b>External</b> . Always unspecified on IRB interfaces.                                                 | <b>detail extensive<br/>brief</b>      |
| <b>Speed</b>            | Speed at which the interface is running. Always unspecified on IRB interfaces.                                                                     | <b>detail extensive<br/>brief</b>      |
| <b>Device flags</b>     | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> . | <b>detail extensive<br/>brief none</b> |
| <b>Interface flags</b>  | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .    | <b>detail extensive<br/>brief none</b> |
| <b>Link type</b>        | Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .                                                                           | <b>detail extensive<br/>none</b>       |
| <b>Link flags</b>       | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .             | <b>detail extensive<br/>none</b>       |
| <b>Physical Info</b>    | Physical interface information.                                                                                                                    | All levels                             |
| <b>Hold-times</b>       | Current interface hold-time up and hold-time down, in milliseconds.                                                                                | <b>detail extensive</b>                |
| <b>Current address</b>  | Configured MAC address.                                                                                                                            | <b>detail extensive<br/>none</b>       |
| <b>Hardware address</b> | MAC address of the hardware.                                                                                                                       | <b>detail extensive<br/>none</b>       |

Table 162: show interfaces irb Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output                 |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Alternate link address</b>  | Backup address of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b>         |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                       | <b>detail extensive</b><br>none |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b>         |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>                                             | <b>detail extensive</b>         |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |

Table 162: show interfaces irb Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output         |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Input errors</b> | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame terminates and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b> |

Table 162: show interfaces irb Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output                 |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Output errors</b>     | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame terminates and FCS errors.</li> <li>• <b>Drops</b>—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.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b>         |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                 |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | All levels                      |
| <b>Index</b>             | Index number of the logical interface (which reflects its initialization sequence).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>      | SNMP interface index number of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b><br>none |
| <b>Generation</b>        | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b>         |

Table 162: show interfaces irb Output Fields (*Continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output         |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Flags</b>              | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> |
| <b>Encapsulation</b>      | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>Bandwidth</b>          | Dummy value that is ignored by an IRB interface. IRB interfaces are pseudo interfaces and do not have physical bandwidth associated with them.                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b> |
| <b>Routing Instance</b>   | Routing instance IRB is configured under.                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b> |
| <b>Bridging Domain</b>    | Bridging domain IRB is participating in.                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail extensive</b> |
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the logical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b> |

Table 162: show interfaces irb Output Fields (*Continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output                 |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |
| <b>Local statistics</b>        | Statistics for traffic received from and transmitted to the Routing Engine.                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b>         |
| <b>Transit statistics</b>      | Statistics for traffic transiting the router.                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>detail extensive</b>         |
| <b>Protocol</b>                | Protocol family configured on the local interface. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>         |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b>         |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b><br>none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b>         |
| <b>Route table</b>             | Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.                                                                                                                                                                                                                                                                                                                                                                             | <b>detail extensive</b>         |

Table 162: show interfaces irb Output Fields (Continued)

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

## Sample Output

### show interfaces irb extensive

```

user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
 Interface index: 129, SNMP ifIndex: 23, Generation: 130
 Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: Unspecified
 Device flags : Present Running
 Interface flags: SNMP-Traps
 Link type : Full-Duplex
 Link flags : None
 Physical info : Unspecified
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
 Alternate link address: Unspecified
 Last flapped : Never
 Statistics last cleared: Never
 Traffic statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0

```

```

Output packets: 0
IPv6 transit statistics:
 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0
Input errors:
 Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed
discards: 0, Resource errors: 0
Output errors:
 Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

```

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)

Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2

Bandwidth: 1000mbps

Routing Instance: customer\_0 Bridging Domain: bd0

Traffic statistics:

```

 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

IPv6 transit statistics:

```

 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

Local statistics:

```

 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

Transit statistics:

```

 Input bytes : 0 0 bps
 Output bytes : 0 0 bps
 Input packets: 0 0 pps
 Output packets: 0 0 pps

```

IPv6 transit statistics:

```

 Input bytes : 0
 Output bytes : 0
 Input packets: 0
 Output packets: 0

```

Protocol inet, MTU: 1500, Generation: 154, Route table: 0

```

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
 Generation: 155
Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
 Flags: Is-Primary
 Policer: Input: __default_arp_policer

```

## show interfaces irb snmp-index

```

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

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
 Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
 Bandwidth: 1000mbps
 Routing Instance: customer_0 Bridging Domain: bd0
 Input packets : 0
 Output packets: 0
 Protocol inet, MTU: 1500
 Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
 Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
 Protocol multiservice, MTU: 1500
 Flags: Is-Primary

```

## show interfaces irb (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM)

```

user@host> show interfaces irb
Physical interface: irb, Enabled, Physical link is Up
 Interface index: 133, SNMP ifIndex: 503

```

```

Type: Ethernet, Link-level type: Ethernet, MTU: 1514
Device flags : Present Running
Interface flags: SNMP-Traps
Link type : Full-Duplex
Link flags : None
Current address: 00:90:69:7f:e7:f0, Hardware address: 00:90:69:7f:e7:f0
Last flapped : Never
 Input packets : 0
 Output packets: 0

Logical interface irb.0 (Index 327) (SNMP ifIndex 507)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
 MAC: 00:11:22:33:44:55
 Bandwidth: 1000mbps
 Routing Instance: default-switch Bridging Domain: bd512+512
 Input packets : 0
 Output packets: 1
 Protocol inet, MTU: 1500
 Flags: Sendbcast-pkt-to-re
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
 Protocol multiservice, MTU: 1500

```

### **show interfaces irb.0 (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM)**

```

user@host> show interfaces irb.0
regress@bloodlust# run show interfaces irb.0
Logical interface irb.0 (Index 327) (SNMP ifIndex 507)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
 MAC: 00:11:22:33:44:55
 Bandwidth: 1000mbps
 Routing Instance: default-switch Bridging Domain: bd512+512
 Input packets : 0
 Output packets: 1
 Protocol inet, MTU: 1500
 Flags: Sendbcast-pkt-to-re
 Addresses, Flags: Is-Preferred Is-Primary
 Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
 Protocol multiservice, MTU: 1500

```

## Release Information

Command introduced in Junos OS Release 8.4.

## show interfaces mac-database

### IN THIS SECTION

- [Syntax | 1680](#)
- [Description | 1680](#)
- [Options | 1680](#)
- [Additional Information | 1681](#)
- [Required Privilege Level | 1681](#)
- [Output Fields | 1681](#)
- [Sample Output | 1685](#)
- [Release Information | 1688](#)

## Syntax

```
show interfaces mac-database (ge-fpc/pic/port | ge-fpc/pic/port.n | aex| xe-
fpc/pic/port | xe-fpc/pic/port.n | et-fpc/pic/port | et-fpc/pic/port.n) <mac-
address mac-address>
```

## Description

(M Series, T Series, MX Series routers, and PTX Series Packet Transport Routers only) Display media access control (MAC) address information for the specified interface.

## Options

***ge-fpc/pic/port***      Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.

|                                |                                                                                                                                         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <i>ge-fpcl picl port.n</i>     | Display MAC addresses that have been learned on a particular logical interface.                                                         |
| <i>aex</i>                     | Display MAC addresses that have been learned on a particular aggregated Ethernet interface.                                             |
| <i>xe-fpcl picl port</i>       | Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.                              |
| <i>xe-fpcl picl port.n</i>     | Display MAC addresses that have been learned on a particular logical interface.                                                         |
| <i>et-fpcl picl port</i>       | Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.                              |
| <i>et-fpcl picl port.n</i>     | Display MAC addresses that have been learned on a particular logical interface.                                                         |
| <i>mac-address mac-address</i> | (Optional) Display detailed MAC address statistics, including policer information for <b>ge</b> , <b>xe</b> , and <b>et</b> interfaces. |

## Additional Information

On IQ2 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.

## Required Privilege Level

view

## Output Fields

[Table 163 on page 1681](#) lists the output fields for the **show interfaces mac-database** command. Output fields are listed in the approximate order in which they appear.

**Table 163: show interfaces mac-database Output Fields**

| Field Name                | Field Description               |
|---------------------------|---------------------------------|
| <b>Physical Interface</b> |                                 |
| <b>Physical interface</b> | Name of the physical interface. |

Table 163: show interfaces mac-database Output Fields (Continued)

| Field Name                      | Field Description                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enabled</b>                  | State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .             |
| <b>Interface index</b>          | Physical interface index number, which reflects its initialization sequence.                                                                              |
| <b>SNMP ifIndex</b>             | SNMP index number for the physical interface.                                                                                                             |
| <b>Description</b>              | Description and name of the interface.                                                                                                                    |
| <b>Link-level type</b>          | Encapsulation being used on the physical interface.                                                                                                       |
| <b>MTU</b>                      | MTU size on the physical interface.                                                                                                                       |
| <b>Speed</b>                    | Speed at which the interface is running.                                                                                                                  |
| <b>Loopback</b>                 | Whether loopback is enabled and the type of loopback: <b>local</b> or <b>remote</b> .                                                                     |
| <b>Source filtering</b>         | Whether source filtering is configured.                                                                                                                   |
| <b>Flow control</b>             | Whether flow control is enabled or disabled.                                                                                                              |
| <b>Minimum links needed</b>     | (Aggregated Ethernet interfaces only) Number of child links that must be operational for the aggregated interface to be operational.                      |
| <b>Minimum bandwidth needed</b> | (Aggregated Ethernet interfaces only) Minimum amount of bandwidth of child links that must be operational for the aggregated interface to be operational. |
| <b>Device flags</b>             | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .        |

Table 163: show interfaces mac-database Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                           |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Current address</b>   | (Aggregated Ethernet interfaces only) Configured MAC address.                                                                                                                                                                                                                                                               |
| <b>Hardware address</b>  | (Aggregated Ethernet interfaces only) Hardware MAC address.                                                                                                                                                                                                                                                                 |
| <b>Last flapped</b>      | (Aggregated Ethernet interfaces only) Date, time, and how long ago the interface went from down to up or from up to down. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (w weeksd days hours:minutes ago)</b> . For example, <b>Last flapped: 2013-12-18 04:33:22 PST (1w5d 22:23 ago)</b> . |
| <b>Input Rate</b>        | (Aggregated Ethernet interfaces only) Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                     |
| <b>Output Rate</b>       | (Aggregated Ethernet interfaces only) Output rate in bps and pps.                                                                                                                                                                                                                                                           |
| <b>Interface flags</b>   | Information about the interface. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                 |
| <b>Link flags</b>        | Information about the link. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                     |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                             |
| <b>Logical interface</b> | Name of the logical interface.                                                                                                                                                                                                                                                                                              |
| <b>Index</b>             | Logical interface index number, which reflects its initialization sequence.                                                                                                                                                                                                                                                 |
| <b>SNMP ifIndex</b>      | Logical interface SNMP interface index number.                                                                                                                                                                                                                                                                              |
| <b>Flags</b>             | Information about the logical interface (possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> ).                                                                                                                                                            |

Table 163: show interfaces mac-database Output Fields (Continued)

| Field Name                                                                                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Encapsulation</b>                                                                       | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>MAC address,<br/>Input frames,<br/>Input bytes,<br/>Output frames,<br/>Output bytes</b> | MAC address and corresponding number of input frames, input bytes, output frames, and output bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Number of MAC addresses</b>                                                             | Number of MAC addresses configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Policer Statistics</b>                                                                  | <p>(Displayed for <b>mac-address</b> option for <b>ge</b>, <b>xe</b>, and <b>et</b> interface types only) Display information about policers applied to a logical interface-MAC pair.</p> <ul style="list-style-type: none"> <li>• <b>Policer type</b>—Type of policer that is out of spec with respect to the configuration. It can be one or more of the following: <ul style="list-style-type: none"> <li>• <b>Input premium</b>—Number of high-priority rating out-of-spec frames or bytes received.</li> <li>• <b>Output premium</b>—Number of high-priority rating out-of-spec frames or bytes sent.</li> <li>• <b>Input aggregate</b>—Total number of out-of-spec frames or bytes received.</li> <li>• <b>Output aggregate</b>—Total number of out-of-spec frames or bytes sent.</li> </ul> </li> <li>• <b>Discarded Frames</b>—Number of discarded frames.</li> <li>• <b>Discarded Bytes</b>—Number of discarded bytes.</li> </ul> |

## Sample Output

### show interfaces mac-database (All MAC Addresses on a Port)

```

user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
 Interface index: 372, SNMP ifIndex: 788
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address Input frames Input bytes Output frames Output bytes
00:00:00:00:00:00 1 56 0 0
00:00:c0:01:01:02 7023810 323095260 0 0
00:00:c0:01:01:03 7023810 323095260 0 0
00:00:c0:01:01:04 7023810 323095260 0 0
00:00:c0:01:01:05 7023810 323095260 0 0
00:00:c0:01:01:06 7023810 323095260 0 0
00:00:c0:01:01:07 7023810 323095260 0 0
00:00:c0:01:01:08 7023809 323095214 0 0
00:00:c0:01:01:09 7023809 323095214 0 0
00:00:c0:01:01:0a 7023809 323095214 0 0
00:00:c0:01:01:0b 7023809 323095214 0 0
00:00:c8:01:01:02 30424784 1399540064 37448598 1722635508
00:00:c8:01:01:03 30424784 1399540064 37448598 1722635508
00:00:c8:01:01:04 30424716 1399536936 37448523 1722632058
00:00:c8:01:01:05 30424789 1399540294 37448598 1722635508
00:00:c8:01:01:06 30424788 1399540248 37448597 1722635462
00:00:c8:01:01:07 30424783 1399540018 37448597 1722635462
00:00:c8:01:01:08 30424783 1399540018 37448596 1722635416
00:00:c8:01:01:09 8836796 406492616 8836795 406492570
00:00:c8:01:01:0a 30424712 1399536752 37448521 1722631966
00:00:c8:01:01:0b 30424715 1399536890 37448523 1722632058
Number of MAC addresses : 21

```

**show interfaces mac-database (All MAC Addresses on an Aggregated Ethernet Interface)**

```

user@host> show interfaces mac-database ae4
Physical interface: ae4, Enabled, Physical link is Up
 Interface index: 132, SNMP ifIndex: 588
 Description: Member links xe-0/2/0
 Link-level type: Ethernet, MTU: 9188, Speed: Unspecified, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled,
 Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
 Device flags : Present Running
 Interface flags: Interface flags: SNMP-Traps Internal: 0x4000
 Current address: 00:22:83:76:ff:c4, Hardware address: 00:22:83:76:ff:c4
 Last flapped : 2013-12-18 04:33:22 PST (1w5d 22:23 ago)
 Input rate : 62756384 bps (85266 pps)
 Output rate : 62759472 bps (85272 pps)

Logical interface ae4.0 (Index 334) (SNMP ifIndex 647)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
 MAC address Input frames Input bytes Output frames Output bytes
 00:00:00:aa:00:02 23888711 2627758118 300 22200
 00:00:00:aa:00:03 0 0 0 0
 00:00:00:aa:00:04 0 0 0 0
 Number of MAC addresses : 3

```

**show interfaces mac-database (All MAC Addresses on a Service)**

```

user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
 MAC address Input frames Input bytes Output frames Output bytes
 00:00:00:00:00:00 1 56 0 0
 00:00:c0:01:01:02 7023810 323095260 0 0
 00:00:c0:01:01:03 7023810 323095260 0 0
 00:00:c0:01:01:04 7023810 323095260 0 0
 00:00:c0:01:01:05 7023810 323095260 0 0
 00:00:c0:01:01:06 7023810 323095260 0 0
 00:00:c0:01:01:07 7023810 323095260 0 0

```

|                   |          |            |          |            |
|-------------------|----------|------------|----------|------------|
| 00:00:c0:01:01:08 | 7023809  | 323095214  | 0        | 0          |
| 00:00:c0:01:01:09 | 7023809  | 323095214  | 0        | 0          |
| 00:00:c0:01:01:0a | 7023809  | 323095214  | 0        | 0          |
| 00:00:c0:01:01:0b | 7023809  | 323095214  | 0        | 0          |
| 00:00:c8:01:01:02 | 31016568 | 1426762128 | 38040381 | 1749857526 |
| 00:00:c8:01:01:03 | 31016568 | 1426762128 | 38040382 | 1749857572 |
| 00:00:c8:01:01:04 | 31016499 | 1426758954 | 38040306 | 1749854076 |
| 00:00:c8:01:01:05 | 31016573 | 1426762358 | 38040381 | 1749857526 |
| 00:00:c8:01:01:06 | 31016573 | 1426762358 | 38040381 | 1749857526 |
| 00:00:c8:01:01:07 | 31016567 | 1426762082 | 38040380 | 1749857480 |
| 00:00:c8:01:01:08 | 31016567 | 1426762082 | 38040379 | 1749857434 |
| 00:00:c8:01:01:09 | 9428580  | 433714680  | 9428580  | 433714680  |
| 00:00:c8:01:01:0a | 31016496 | 1426758816 | 38040304 | 1749853984 |
| 00:00:c8:01:01:0b | 31016498 | 1426758908 | 38040307 | 1749854122 |

### show interfaces mac-database mac-address

```

user@host> show interfaces mac-database xe-0/3/3 mac-address 00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
 Interface index: 372, SNMP ifIndex: 788
 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
 Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address: 00:00:c8:01:01:09, Type: Configured,
 Input bytes : 202324652
 Output bytes : 202324560
 Input frames : 4398362
 Output frames : 4398360
Policer statistics:
Policer type Discarded frames Discarded bytes
Output aggregate 3992386 183649756

```

## Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

Support for statement with the **aeX** option introduced in Junos OS Release 15.1.

## show interfaces mc-ae

### IN THIS SECTION

- [Syntax | 1688](#)
- [Description | 1688](#)
- [Options | 1689](#)
- [Required Privilege Level | 1689](#)
- [Output Fields | 1689](#)
- [Sample Output | 1692](#)
- [Release Information | 1695](#)

## Syntax

```
show interfaces mc-ae
 extensive
 revertive-info
 <id identifier unit number>
```

## Description

On MX Series routers with multichassis aggregated Ethernet (**aeX**) interfaces, displays information about the **aeX** interfaces.

## Options

- extensive** (Optional) Display extensive information for multichassis aggregated Ethernet interface.
- revertive-info** (Optional) Display revertive mode information for multichassis aggregated Ethernet interface.
- identifier** (Optional) Identifier of the multichassis aggregated Ethernet interface.
- number** (Optional) Specify the logical interface by unit number.

## Required Privilege Level

view

## Output Fields

[Table 164 on page 1689](#) lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

**Table 164: show interfaces mc-ae Output Fields**

| Output Field Name                    | Field Description                                                                                                              |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Member Link</b>                   | Identifiers of the configured multichassis link aggregate interfaces configured interfaces.                                    |
| <b>Current State Machine's State</b> | Current state of the MCLAG state machine. The MCLAG state machine is responsible for synchronization with the peer MCLAG node. |
| <b>Local Status</b>                  | Status of the local link: <b>active</b> or <b>standby</b> .                                                                    |
| <b>Peer Status</b>                   | Status of the peer link: <b>active</b> or <b>standby</b> .                                                                     |
| <b>Local State</b>                   | Up or down state of the local device.                                                                                          |

Table 164: show interfaces mc-ae Output Fields (Continued)

| Output Field Name                 | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Peer State</b>                 | Status of the local and peer links in an <b>active/active</b> bridge or VRRP over integrated routing and bridging (IRB) configuration on MX Series routers, including: <ul style="list-style-type: none"> <li><b>Logical Interface</b>      Aggregated Ethernet (AE) aggregate number and unit number.</li> <li><b>Topology Type</b>        The bridge or VRRP topology type configured on the AE.</li> <li><b>Local State</b>            Up or down state of the local device.</li> <li><b>Peer State</b>              Up or down state of the peer device.</li> <li><b>Peer Ip/ICL-PL/State</b>    Address, interface and state of the peer device.</li> </ul> |
| <b>Core Facing Interface</b>      | Label: <b>pseudowire interface</b> or <b>Ethernet interface</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>ICL-PL</b>                     | Label: <b>pseudowire interface</b> or <b>Ethernet interface</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>switchover mode</b>            | The configured switchover mode for the multichassis aggregated Ethernet interface: <b>revertive</b> or <b>non-revertive</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>switchover status</b>          | Status of the switchover if the <b>revert-time</b> statement is configured at the <b>[edit interfaces aexmc-ae]</b> hierarchy level.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>revert time</b>                | Revert time configured for the multichassis aggregated Ethernet interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>switchover time remaining</b>  | Seconds left to trigger the switchover if the switchover is in progress.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Configuration Error Status</b> | Reason for the configuration error.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

Table 164: show interfaces mc-ae Output Fields (Continued)

| Output Field Name        | Field Description                                                                                                                                                                                                                        |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MCAE Configuration       |                                                                                                                                                                                                                                          |
| <b>Redundancy Groups</b> | Identification number of the redundancy group. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate multiple chassis that perform similar redundancy functions. Possible values: 1 through 4,294,967,294. |
| <b>MCAE ID</b>           | Identification number of the MC-LAG device. The two MC-LAG network devices that manage a given MC-LAG must have the same identification number. Possible values:1 through 65,535.                                                        |
| <b>MCAE Mode</b>         | Specifies whether the MC-LAG is in active-active or active-standby mode. Possible values: active-active or active-standby.                                                                                                               |
| <b>Status Control</b>    | Specifies whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs. Possible values: active or standby.                                                                                    |
| <b>Chassis ID</b>        | Chassis ID for Link Aggregation Control Protocol (LACP) to calculate the port number of MC-LAG physical member links. Possible values: 0 or 1.                                                                                           |
| LACP Configuration       |                                                                                                                                                                                                                                          |
| <b>System ID</b>         | System id of the local system.                                                                                                                                                                                                           |
| <b>Admin Key</b>         | LACP administrative key of the node. Value should be the same on both MCLAG peers.                                                                                                                                                       |
| LACP Information         |                                                                                                                                                                                                                                          |

**Table 164: show interfaces mc-ae Output Fields (Continued)**

| Output Field Name              | Field Description                                  |
|--------------------------------|----------------------------------------------------|
| <b>Local Partner System ID</b> | LACP partner system ID as seen by the local node.  |
| <b>Peer Partner System ID</b>  | LACP partner system ID as seen by MC-AE peer node. |

## Sample Output

### show interfaces mc-ae

```

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

```

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

```

user@host# show interfaces mc-ae ge-0/0/0.0
Member Link : ae0
Current State Machine's State: active
Local Status : active
Local State : up
Peer Status : active
Peer State : up
 Logical Interface : ae0.0
 Topology Type : bridge
 Local State : up
 Peer State : up
 Peer Ip/ICL-PL/State : 192.168.100.10 ge-0/0/0.0 up

```

**show interfaces mc-ae revertive-info**

```

user@host> show interfaces mc-ae revertive-info id 2
Member Link : ae1
Current State Machine's State: mcae active state
Local Status : active
Local State : up
Peer Status : standby
Peer State : up
Switchover Mode : Non Revertive
Switchover Status : N/A
Revert Time : 1 Minutes
Switchover Remaining Time : N/A
 Logical Interface : ae1.1024
 Topology Type : bridge
 Local State : up
 Peer State : up
 Peer Ip/MCP/State : N/A

```

**show interfaces mc-ae extensive**

```

user@host> show interfaces mc-ae extensive
Member Link : ae2
Current State Machine's State: mcae active state
Local Status : active
Local State : up
Peer Status : active
Peer State : up
 Logical Interface : ae2.1
 Topology Type : bridge
 Local State : up
 Peer State : up
 Peer Ip/MCP/State : 192.168.143.17 ae0.1 up

MCAE Configuration
 Redundancy Group : 1
 MCAE ID : 2
 MCAE Mode : active_active
 Status Control : active
 Chassis ID : 0

```

```
LACP Configuration
 System ID : 00:00:00:00:00:02
 Admin Key : 10
```

### show interfaces mc-ae extensive (MX Series Router after a configuration exchange error)

```
user@host> show interfaces mc-ae extensive
Member Link : ae2
Current State Machine's State: mcae config exchange error
Configuration Error Status : same chassis-id
Local Status : active
Local State : up
Peer Status : Unknown
Peer State : Unknown
 Logical Interface : ae2.1
 Topology Type : bridge
 Local State : up
 Peer State : up
 Peer Ip/MCP/State : 192.168.143.17 ae0.1 up

MCAE Configuration
 Redundancy Group : 1
 MCAE ID : 2
 MCAE Mode : active_active
 Status Control : active
 Chassis ID : 1

LACP Configuration
 System ID : 00:00:00:00:00:02
 Admin Key : 10
```

### show interfaces mc-ae extensive

```
user@host> show interfaces mc-ae extensive
Member Link : ae0
Current State Machine's State: mcae active state
Local Status : active
Local State : up
Peer Status : active
Peer State : up
```

```

Logical Interface : ae0.1
Topology Type : bridge
Local State : up
Peer State : up
Peer Ip/MCP/State : 192.168.143.17 ge-0/0/2.1 up

MCAE Configuration
Redundancy Group : 1
MCAE ID : 1
MCAE Mode : active_active
Status Control : active
Chassis ID : 0

LACP Configuration
System ID : 00:00:00:00:00:02
Admin Key : 10

LACP Information
Local Partner System ID : 00:00:00:00:00:01
Peer Partner System ID : 00:00:00:00:00:01

```

## Release Information

Command introduced in Junos OS Release 9.6.

**revertive-info** statement introduced in Junos OS Release 13.3

**extensive** statement introduced in Junos OS Release 16.1R1

## RELATED DOCUMENTATION

| [Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up](#)

## show interfaces port-profile

### IN THIS SECTION

- [Syntax | 1696](#)
- [Description | 1696](#)

- Options | [1696](#)
- Required Privilege Level | [1696](#)
- Output Fields | [1696](#)
- Sample Output | [1697](#)
- Release Information | [1703](#)

## Syntax

```
show interfaces port-profile
```

## Description

Display information about the operating port speed summary for the line card. You can configure channelized and non-channelized interfaces on each physical port on a PIC. The different channelized and non-channelized interfaces can operate at different speeds. Use the command to view the port speeds for the interfaces (channelized and non-channelized) configured on the line card.

## Options

This command has no options.

## Required Privilege Level

view

## Output Fields

[Table 165 on page 1697](#) lists the output fields for the **show interfaces port-profile** commands. Output fields are listed in the approximate order in which they appear.

**Table 165: show interfaces port-profile Output fields**

| Field Name | Description                                                                      |
|------------|----------------------------------------------------------------------------------|
| Interface  | Name of the channelized or non-channelized interface.                            |
| Port Speed | Speed configured on the port for the interface (channelized or non-channelized). |

## Sample Output

### show interfaces port-profile (PTX10001-36MR)

```
user@host> show interfaces port-profile
```

```
Interface Port Speed
 et-0/0/0 400Gbps
 et-0/0/1:0 25Gbps
 et-0/0/1:1 25Gbps
 et-0/0/1:2 25Gbps
 et-0/0/1:3 25Gbps
 et-0/0/1:4 25Gbps
 et-0/0/1:5 25Gbps
 et-0/0/1:6 25Gbps
 et-0/0/1:7 25Gbps
 et-0/0/2 40Gbps
 et-0/0/3:0 10Gbps
 et-0/0/3:1 10Gbps
 et-0/0/3:2 10Gbps
 et-0/0/3:3 10Gbps
 et-0/0/4 100Gbps
 et-0/0/5 100Gbps
 et-0/0/6 100Gbps
 et-0/0/7 100Gbps
 et-0/0/8:0 100Gbps
 et-0/0/8:1 100Gbps
 et-0/0/9 400Gbps
 et-0/0/10:0 50Gbps
 et-0/0/10:1 50Gbps
 et-0/0/10:2 50Gbps
 et-0/0/10:3 50Gbps
```

```
et-0/0/10:4 50Gbps
et-0/0/10:5 50Gbps
et-0/0/10:6 50Gbps
et-0/0/10:7 50Gbps
et-0/0/11 400Gbps
et-0/1/0 400Gbps
et-0/1/1:0 100Gbps
et-0/1/1:1 100Gbps
et-0/1/2 400Gbps
et-0/1/3 100Gbps
et-0/1/4 100Gbps
```

### show interfaces port-profile (PTX10008 Series routers with JNP10K-LC1201)

```
user@host> show interfaces port-profile
```

```
Interface Port Speed
et-7/0/0:0 100Gbps
et-7/0/0:1 100Gbps
et-7/0/0:2 100Gbps
et-7/0/0:3 100Gbps
et-7/0/1 100Gbps
et-7/0/2 400Gbps
et-7/0/3 400Gbps
et-7/0/4 400Gbps
et-7/0/5 400Gbps
et-7/0/6 400Gbps
et-7/0/7 400Gbps
et-7/0/8 400Gbps
et-7/0/9 400Gbps
et-7/0/10 400Gbps
et-7/0/11 400Gbps
et-7/0/12 400Gbps
et-7/0/13 400Gbps
et-7/0/14 400Gbps
et-7/0/15 400Gbps
et-7/0/16 400Gbps
et-7/0/17 400Gbps
et-7/0/18 400Gbps
et-7/0/19 400Gbps
et-7/0/20 400Gbps
et-7/0/21 400Gbps
```

```
et-7/0/22 400Gbps
et-7/0/23 400Gbps
et-7/0/24 400Gbps
et-7/0/25 400Gbps
et-7/0/26 400Gbps
et-7/0/27 400Gbps
et-7/0/28 400Gbps
et-7/0/29 400Gbps
et-7/0/30 400Gbps
et-7/0/31 400Gbps
et-7/0/32 400Gbps
et-7/0/33 400Gbps
et-7/0/34 400Gbps
et-7/0/35 400Gbps
```

### show interfaces port-profile (PTX10008 Series routers with JNP10K-LC1201)

```
user@host> show interfaces port-profile
Interface Port Speed
et-0/0/1 400Gbps
et-0/0/2 400Gbps
et-0/0/3 400Gbps
et-0/0/4 400Gbps
et-0/0/5 400Gbps
et-0/0/6 400Gbps
et-0/0/7 400Gbps
et-0/0/8 400Gbps
et-0/0/9 400Gbps
et-0/0/10 400Gbps
et-0/0/11 400Gbps
et-0/0/12 400Gbps
et-0/0/13 400Gbps
et-0/0/14 400Gbps
et-0/0/15 400Gbps
et-0/0/16 400Gbps
et-0/0/17 10Gbps
et-0/0/18 400Gbps
et-0/0/19 400Gbps
et-0/0/20 400Gbps
et-0/0/21 400Gbps
et-0/0/22 400Gbps
```

|           |         |
|-----------|---------|
| et-0/0/23 | 400Gbps |
| et-0/0/24 | 400Gbps |
| et-0/0/25 | 400Gbps |
| et-0/0/26 | 400Gbps |
| et-0/0/27 | 400Gbps |
| et-0/0/28 | 400Gbps |
| et-0/0/29 | 400Gbps |
| et-0/0/30 | 400Gbps |
| et-0/0/31 | 400Gbps |
| et-0/0/32 | 400Gbps |
| et-0/0/33 | 400Gbps |
| et-0/0/34 | 400Gbps |
| et-0/0/35 | 400Gbps |
| et-3/0/0  | 400Gbps |
| et-3/0/1  | 400Gbps |
| et-3/0/2  | 400Gbps |
| et-3/0/3  | 400Gbps |
| et-3/0/4  | 400Gbps |
| et-3/0/5  | 100Gbps |
| et-3/0/6  | 400Gbps |
| et-3/0/7  | 400Gbps |
| et-3/0/8  | 400Gbps |
| et-3/0/9  | 400Gbps |
| et-3/0/10 | 400Gbps |
| et-3/0/11 | 400Gbps |
| et-3/0/12 | 400Gbps |
| et-3/0/13 | 400Gbps |
| et-3/0/14 | 400Gbps |
| et-3/0/15 | 400Gbps |
| et-3/0/16 | 400Gbps |
| et-3/0/17 | 400Gbps |
| et-3/0/18 | 400Gbps |
| et-3/0/19 | 400Gbps |
| et-3/0/20 | 400Gbps |
| et-3/0/21 | 400Gbps |
| et-3/0/22 | 400Gbps |
| et-3/0/23 | 400Gbps |
| et-3/0/24 | 400Gbps |
| et-3/0/25 | 400Gbps |
| et-3/0/26 | 400Gbps |
| et-3/0/27 | 400Gbps |
| et-3/0/28 | 400Gbps |
| et-3/0/29 | 400Gbps |

|             |         |
|-------------|---------|
| et-3/0/30   | 400Gbps |
| et-3/0/31   | 10Gbps  |
| et-3/0/32   | 10Gbps  |
| et-3/0/33   | 400Gbps |
| et-3/0/34   | 400Gbps |
| et-3/0/35   | 400Gbps |
| et-4/0/0    | 400Gbps |
| et-4/0/1    | 100Gbps |
| et-4/0/2    | 100Gbps |
| et-4/0/3    | 100Gbps |
| et-4/0/4    | 400Gbps |
| et-4/0/5    | 400Gbps |
| et-4/0/6    | 100Gbps |
| et-4/0/7    | 100Gbps |
| et-4/0/8:0  | 100Gbps |
| et-4/0/8:1  | 100Gbps |
| et-4/0/9    | 400Gbps |
| et-4/0/10   | 100Gbps |
| et-4/0/11   | 40Gbps  |
| et-4/0/12   | 100Gbps |
| et-4/0/13   | 400Gbps |
| et-4/0/14   | 400Gbps |
| et-4/0/15   | 400Gbps |
| et-4/0/16:0 | 100Gbps |
| et-4/0/16:1 | 100Gbps |
| et-4/0/17   | 400Gbps |
| et-4/0/18   | 400Gbps |
| et-4/0/19   | 400Gbps |
| et-4/0/20   | 400Gbps |
| et-4/0/21   | 400Gbps |
| et-4/0/22   | 400Gbps |
| et-4/0/23   | 100Gbps |
| et-4/0/24   | 400Gbps |
| et-4/0/25   | 400Gbps |
| et-4/0/26   | 400Gbps |
| et-4/0/27   | 400Gbps |
| et-4/0/28   | 400Gbps |
| et-4/0/29   | 400Gbps |
| et-4/0/30   | 400Gbps |
| et-4/0/31   | 100Gbps |
| et-4/0/32   | 100Gbps |
| et-4/0/33   | 400Gbps |
| et-4/0/34   | 400Gbps |

|             |         |
|-------------|---------|
| et-4/0/35   | 400Gbps |
| et-5/0/0    | 400Gbps |
| et-5/0/1:0  | 100Gbps |
| et-5/0/1:1  | 100Gbps |
| et-5/0/2    | 100Gbps |
| et-5/0/3    | 100Gbps |
| et-5/0/4    | 400Gbps |
| et-5/0/5    | 40Gbps  |
| et-5/0/6:0  | 25Gbps  |
| et-5/0/6:1  | 25Gbps  |
| et-5/0/6:2  | 25Gbps  |
| et-5/0/6:3  | 25Gbps  |
| et-5/0/6:4  | 25Gbps  |
| et-5/0/6:5  | 25Gbps  |
| et-5/0/6:6  | 25Gbps  |
| et-5/0/6:7  | 25Gbps  |
| et-5/0/7:0  | 10Gbps  |
| et-5/0/7:1  | 10Gbps  |
| et-5/0/7:2  | 10Gbps  |
| et-5/0/7:3  | 10Gbps  |
| et-5/0/8:0  | 100Gbps |
| et-5/0/8:1  | 100Gbps |
| et-5/0/9    | 40Gbps  |
| et-5/0/10   | 100Gbps |
| et-5/0/11   | 400Gbps |
| et-5/0/12   | 100Gbps |
| et-5/0/13:0 | 10Gbps  |
| et-5/0/13:1 | 10Gbps  |
| et-5/0/13:2 | 10Gbps  |
| et-5/0/13:3 | 10Gbps  |
| et-5/0/14   | 400Gbps |
| et-5/0/15   | 40Gbps  |
| et-5/0/16   | 400Gbps |
| et-5/0/17   | 400Gbps |
| et-5/0/18   | 400Gbps |
| et-5/0/19   | 400Gbps |
| et-5/0/20   | 400Gbps |
| et-5/0/21   | 400Gbps |
| et-5/0/22   | 400Gbps |
| et-5/0/23   | 400Gbps |
| et-5/0/24   | 400Gbps |
| et-5/0/25   | 400Gbps |
| et-5/0/26   | 400Gbps |

```
et-5/0/27 400Gbps
et-5/0/28 400Gbps
et-5/0/29 400Gbps
et-5/0/30 400Gbps
et-5/0/31 400Gbps
et-5/0/32 100Gbps
et-5/0/33 100Gbps
et-5/0/34:0 10Gbps
et-5/0/34:1 10Gbps
et-5/0/34:2 10Gbps
et-5/0/34:3 10Gbps
et-5/0/35 400Gbps
```

## Release Information

Command introduced in Junos OS Evolved Release 20.1R2.

## show interfaces prbs-stats

### IN THIS SECTION

- [Syntax \(MX10003 and MX204\) | 1703](#)
- [Description | 1704](#)
- [Required Privilege Level | 1704](#)
- [Sample Output | 1704](#)
- [Release Information | 1704](#)

## Syntax (MX10003 and MX204)

```
show interfaces interface-name prbs-stats
```

## Description

Displays the Pseudo Random Binary Sequence (PRBS) statistics and the status of the test (PASS/FAIL) along with error counters.

Use the "[prbs-test-start](#)" on page 1239 and "[prbs-test-stop](#)" on page 1242 commands to run and stop the PRBS statistics collection respectively.

For the step-by-step procedure on how to collect and view the PRBS statistics, refer "[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#)" on page 681.

A 10-Gigabit ethernet interface contains a single lane of data transmission and reception, while a 40-Gigabit ethernet and 100-Gigabit ethernet interface comprises of four lanes of 10G and 25G for data transmission and reception respectively. The PRBS tests are executed per lane of an interface. Hence, The PRBS test status displays the status per lane of the interface.

## Required Privilege Level

view

## Sample Output

**show interfaces et-0/1/2 prbs-stats (MX10003 and MX204 routers)**

```
user@host> 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 tests are executed per lane of an interface. Hence, The PRBS test status displays the status per lane of the interface.

## Release Information

Statement introduced in Junos OS Release 19.2R1.

## RELATED DOCUMENTATION

[prbs-test-start](#) | [1239](#)

[prbs-test-stop](#) | [1242](#)

*clear interfaces statistics*

[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#) | [681](#)

## show interfaces smart-sfp-defects

### IN THIS SECTION

- [Syntax](#) | [1705](#)
- [Description](#) | [1705](#)
- [Options](#) | [1706](#)
- [Required Privilege Level](#) | [1706](#)
- [Output Fields](#) | [1706](#)
- [Sample Output](#) | [1716](#)
- [Release Information](#) | [1719](#)

### Syntax

```
show interfaces smart-sfp-defects
<brief | detail | terse>
<interface-name>
```

### Description

Display information about the defects on the smart SFP interface. The defects can be Smart SFP device defects or TDM legacy defects.

## Options

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

***interface-name*** Name of the interface.

## Required Privilege Level

view

## Output Fields

[Table 166 on page 1706](#) lists the output fields for the **show interfaces smart-sfp-defects** command. Output fields are listed in the approximate order in which they appear.

**Table 166: show interfaces smart-sfp-defects Output Fields**

| Field Name                | Field Description                                                                   | Level of Output |
|---------------------------|-------------------------------------------------------------------------------------|-----------------|
| <b>Physical Interface</b> |                                                                                     |                 |
| Physical interface        | Name of the physical interface.                                                     | All levels      |
| Enabled                   | State of the interface.                                                             | All levels      |
| Interface index           | Index number of the physical interface, which reflects its initialization sequence. | detail none     |
| SNMP ifIndex              | SNMP index number for the physical interface.                                       | detail none     |
| Link-level type           | Encapsulation being used on the physical interface.                                 | All levels      |
| MTU                       | Maximum transmission unit size on the physical interface.                           | All levels      |
| Clocking                  | Reference clock source.                                                             | detail          |

**Table 166: show interfaces smart-sfp-defects Output Fields (Continued)**

| Field Name             | Field Description                                                                                                                 | Level of Output |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Speed                  | Speed at which the interface is running.                                                                                          | All levels      |
| Duplex                 | Duplex mode of the interface, either Full-Duplex or Half-Duplex.                                                                  | All levels      |
| MAC-REWRITE Error      | Specifies if the encapsulation of the packet has been changed.                                                                    | none            |
| BPDU Error             | Specifies if a BPDU has been received on a blocked interface.                                                                     | none            |
| Loopback               | Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.                                  | All levels      |
| Source filtering       | Source filtering status: Enabled or Disabled.                                                                                     | All levels      |
| Flow control           | Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.              | All levels      |
| Device flags           | Information about the physical device.                                                                                            | All levels      |
| Interface flags        | Information about the interface.                                                                                                  | All levels      |
| Smart Transceiver Type | Type of smart SFP transceiver. Possible values: E1, T1, DS3, STM1, STM4, or STM16.                                                | All levels      |
| Smart SFP              | Configurable Type. You can configure these Smart SFP transceivers unlike other Smart SFP transceivers which you cannot configure. | All levels      |

**Table 166: show interfaces smart-sfp-defects Output Fields (Continued)**

| Field Name               | Field Description                                                                   | Level of Output |
|--------------------------|-------------------------------------------------------------------------------------|-----------------|
| Smart SFP Configurations |                                                                                     |                 |
| Loopback                 | Specifies if loopback is configured.                                                | All levels      |
| Encapsulation Circuit Id | Emulation circuit id for encapsulation. Possible values: 0 through 1048575.         | All levels      |
| Decapsulation Circuit Id | Emulation circuit id for decapsulation. Possible values: 0 through 1048575.         | All levels      |
| Mode                     | Framing mode. Possible values: MEF8 or MPLS.                                        | All levels      |
| Dest MAC                 | Destination MAC Address                                                             | detail          |
| Smart SFP Defects        |                                                                                     |                 |
| TDM defects              | Legacy TDM defects. Defects specific to PDH and SDH legacy error codes and defects. | All             |
| System defects           | Smart SFP device defects. Defects are specific to the Smart SFP you use.            | All             |
| CoS queues               | Number of CoS queues configured.                                                    | detail none     |
| Hold-Times               | Current interface hold-time up and hold-time down, in milliseconds.                 | detail          |
| Current address          | Configured MAC address.                                                             | detail none     |

Table 166: show interfaces smart-sfp-defects Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Hardware address        | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | detail none     |
| Last flapped            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .                                                                                                                                                                                                                                    | detail none     |
| Statistics last cleared | Date, time, and how long ago the statistics for the interface were cleared. The format is <b>Statistics last cleared: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>2010-05-17 07:51:28 PDT (00:04:33 ago)</b> .                                                                                                                                                                                                                            | detail          |
| Traffic statistics      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |

Table 166: show interfaces smart-sfp-defects Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| IPv6 transit statistics | <p>If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface:</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |

Table 166: show interfaces smart-sfp-defects Output Fields (Continued)

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Input errors | <p data-bbox="431 369 1159 436">Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul data-bbox="431 474 1187 1812" style="list-style-type: none"> <li data-bbox="431 474 1110 541">• Errors—Sum of the incoming frame terminates and FCS errors.</li> <li data-bbox="431 579 1175 726">• Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li data-bbox="431 764 1162 831">• Framing errors—Number of packets received with an invalid frame checksum (FCS).</li> <li data-bbox="431 869 1167 936">• Runts—Number of frames received that are smaller than the runt threshold.</li> <li data-bbox="431 974 1182 1121">• Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li data-bbox="431 1159 1166 1348">• L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement.</li> <li data-bbox="431 1386 1183 1453">• L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li data-bbox="431 1491 1159 1600">• L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li data-bbox="431 1638 1170 1747">• FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li data-bbox="431 1785 938 1812">• Resource errors—Sum of transmit drops.</li> </ul> | detail none     |

Table 166: show interfaces smart-sfp-defects Output Fields (Continued)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Output errors | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• Errors—Sum of the outgoing frame terminates and FCS errors.</li> <li>• Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces.</li> </ul> | detail none     |

Table 166: show interfaces smart-sfp-defects Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                              | Level of Output |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|                          | <ul style="list-style-type: none"> <li>• MTU errors—Number of packets whose size exceeded the MTU of the interface.</li> <li>• Resource errors—Sum of transmit drops.</li> </ul>                                                                                                                                               |                 |
| Egress queues            | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                            | detail          |
| Queue counters           | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• Queued packets—Number of queued packets.</li> <li>• Transmitted packets—Number of transmitted packets.</li> <li>• Dropped packets—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> | detail          |
| Input rate               | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                              | None specified  |
| Output rate              | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                    | None specified  |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                                |                 |
| Logical interface        | Name of the logical interface.                                                                                                                                                                                                                                                                                                 | All levels      |
| Index                    | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                             | detail none     |
| SNMP ifIndex             | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                         | detail none     |

Table 166: show interfaces smart-sfp-defects Output Fields (Continued)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Flags              | <p>Information about the logical interface.</p> <p>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.</p>                                                                                      | All levels      |
| Input packets      | Number of packets received on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                   | detail none     |
| Output packets     | Number of packets transmitted on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail none     |
| Encapsulation      | <p>Encapsulation method used on the logical interface.</p> <ul style="list-style-type: none"> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                            | All levels      |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |

Table 166: show interfaces smart-sfp-defects Output Fields (Continued)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Local statistics       | <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |
| Transit statistics     | <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |
| Addresses, Flags       | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                           | detail none     |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                       | brief           |
| MTU                    | Maximum transmission unit size on the physical interface.                                                                                                                                                                                                                                                                                                                                      | All levels      |
| Destination            | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                               | detail none     |
| Local                  | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                           | detail none     |

**Table 166: show interfaces smart-sfp-defects Output Fields (Continued)**

| Field Name  | Field Description                                                                                                 | Level of Output |
|-------------|-------------------------------------------------------------------------------------------------------------------|-----------------|
| Broadcast   | Broadcast address of the logical interlace.                                                                       | detail none     |
| Generation  | Unique number for use by Juniper Networks technical support only.                                                 | detail          |
| Route table | Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0. | detail none     |

## Sample Output

### show interfaces smart-sfp-defects (T1)

```

user@host> show interfaces ge-4/1/0 smart-sfp-defects
Physical interface: ge-4/1/0, Enabled, Physical link is Up
 Interface index: 216, SNMP ifIndex: 742
 Link-level type: Ethernet-Bridge, 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: 0x20004000
 Link flags : None
 Smart Transceiver Type: T1
 Smart SFP: Configurable Type
 Smart SFP Configurations
 Loopback: None
 DMAC check Enabled: enabled
 Encapsulation Circuit Id: 400
 Decapsulation Circuit Id: 16
 VLAN1: 400 VLAN2: 400
 Mode: MEF8

```

```

Dest MAC: 00:19:3a:02:a4:b8
Smart SFP Defects:
 TDM defects : Local Packet Loss
 System defects : None

Logical interface ge-4/1/0.0 (Index 395) (SNMP ifIndex 615)
 Flags: Up SNMP-Traps 0x24004000 Encapsulation: Ethernet-Bridge
 Tenant Name: (null)
 Input packets : 2490261
 Output packets: 2486756
 Protocol bridge, MTU: 1514

```

### show interfaces smart-sfp-defects (DS3)

```

user@host> show interfaces ge-4/0/5 smart-sfp-defects
Physical interface: ge-4/0/5, Enabled, Physical link is Up
 Interface index: 211, SNMP ifIndex: 735
 Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed:
1000mbps, BPDU Error: None, Loop Detect PDU Error: None,
 Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled,
 Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
 Pad to minimum frame size: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None
 Smart Transceiver Type: DS3
 Smart SFP: Configurable Type
 Smart SFP Configurations
 Loopback: None
 Encapsulation Circuit Id: 16
 Decapsulation Circuit Id: 16
 Mode: MPLS
 Smart SFP Defects:
 TDM defects : Loss of Signal, Local Packet Loss, Out Of Frame
 System defects : None

Logical interface ge-4/0/5.0 (Index 355) (SNMP ifIndex 612)
 Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
 Input packets : 430

```

```
Output packets: 0
Protocol multiservice, MTU: Unlimited
```

### show interfaces smart-sfp-defects (STM1)

```
user@host > show interfaces ge-4/0/0 smart-sfp-defects
Physical interface: ge-4/0/0, Enabled, Physical link is Up
 Interface index: 192, SNMP ifIndex: 546, Generation: 269
 Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed:
1000mbps, BPDU Error: None,
 Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE
Error: None,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-
negotiation: Enabled,
 Remote fault: Online
 Pad to minimum frame size: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x4000
 Link flags : None
 Smart Transceiver Type: STM1
 Smart SFP: Configurable Type
 Smart SFP Configurations
 Loopback: None
 DMAC check Enabled: enabled
 Encapsulation Circuit Id: 1000
 Decapsulation Circuit Id: 2000
 VLAN1: 100
 Mode: MEF8
 Dest MAC: 10:0e:7e:37:cd:29
 Smart SFP Defects:
 TDM defects : Local Packet Loss
 System defects : None

Logical interface ge-4/0/0.0 (Index 369) (SNMP ifIndex 799) (Generation 190)
 Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
 Input packets : 2400
 Output packets: 0
 Protocol multiservice, MTU: Unlimited
```

## Release Information

Command introduced in Junos OS Release 19.4.

### RELATED DOCUMENTATION

[Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | 246](#)

[ces-psn-channel \(tdm-options\) | 823](#)

[iwf-params \(tdm-options\) | 934](#)

[tdm-options \(Interfaces\) | 1176](#)

## show interfaces smart-sfp-statistics

### IN THIS SECTION

- [Syntax | 1719](#)
- [Description | 1720](#)
- [Options | 1720](#)
- [Required Privilege Level | 1720](#)
- [Output Fields | 1720](#)
- [Sample Output | 1730](#)
- [Release Information | 1736](#)

## Syntax

```
show interfaces smart-sfp-statistics
<brief | detail | extensive | terse>
<descriptions>
<interface-name>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
```

## Description

Display status information about the specified smart SFP interface.

## Options

|                                                      |                                                                                         |
|------------------------------------------------------|-----------------------------------------------------------------------------------------|
| <b>brief   detail   extensive   terse</b>            | (Optional) Display the specified level of output.                                       |
| <b>descriptions</b>                                  | (Optional) Display interface description strings.                                       |
| <b><i>interface-name</i></b>                         | Name of the interface.                                                                  |
| <b>media</b>                                         | (Optional) Display media-specific information about network interfaces.                 |
| <b>routing-instance (all   <i>instance-name</i>)</b> | (Optional) Display all routing instances or the name of an individual routing instance. |
| <b>snmp-index <i>snmp-index</i></b>                  | (Optional) Display information for the specified SNMP index of the interface.           |

## Required Privilege Level

view

## Output Fields

[Table 167 on page 1720](#) lists the output fields for the **show interfaces smart-sfp-statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 167: show interfaces smart-sfp-statistics Output Fields**

| Field Name                | Field Description               | Level of Output |
|---------------------------|---------------------------------|-----------------|
| <b>Physical Interface</b> |                                 |                 |
| Physical interface        | Name of the physical interface. | All levels      |
| Enabled                   | State of the interface.         | All levels      |

**Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)**

| Field Name        | Field Description                                                                                                    | Level of Output |
|-------------------|----------------------------------------------------------------------------------------------------------------------|-----------------|
| Interface index   | Index number of the physical interface, which reflects its initialization sequence.                                  | detail none     |
| SNMP ifIndex      | SNMP index number for the physical interface.                                                                        | detail none     |
| Link-level type   | Encapsulation being used on the physical interface.                                                                  | All levels      |
| MTU               | Maximum transmission unit size on the physical interface.                                                            | All levels      |
| Clocking          | Reference clock source.                                                                                              | detail          |
| Speed             | Speed at which the interface is running.                                                                             | All levels      |
| Duplex            | Duplex mode of the interface, either Full-Duplex or Half-Duplex.                                                     | All levels      |
| MAC-REWRITE Error | Specifies if the encapsulation of the packet has been changed.                                                       | none            |
| BPDU Error        | Specifies if a BPDU has been received on a blocked interface.                                                        | none            |
| Loopback          | Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.                     | All levels      |
| Source filtering  | Source filtering status: Enabled or Disabled.                                                                        | All levels      |
| Flow control      | Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured. | All levels      |

**Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)**

| Field Name               | Field Description                                                                                                                 | Level of Output |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Device flags             | Information about the physical device.                                                                                            | All levels      |
| Interface flags          | Information about the interface.                                                                                                  | All levels      |
| Smart Transceiver Type   | Type of smart SFP transceiver. Possible values: E1, T1, DS3, STM1, STM4, or STM16.                                                | All levels      |
| Smart SFP                | Configurable Type. You can configure these Smart SFP transceivers unlike other Smart SFP transceivers which you cannot configure. | All levels      |
| Smart SFP Configurations |                                                                                                                                   |                 |
| Loopback                 | Specifies if loopback is configured.                                                                                              | All levels      |
| Encapsulation Circuit Id | Emulation circuit id for encapsulation. Possible values: 0 through 1048575.                                                       | all             |
| Decapsulation Circuit Id | Emulation circuit id for decapsulation. Possible values: 0 through 1048575.                                                       | all             |
| Mode                     | Framing mode. Possible values: MEF8 or MPLS.                                                                                      | all             |
| Dest MAC                 | Destination MAC address.                                                                                                          | all             |
| Smart SFP Defects        |                                                                                                                                   |                 |
| TDM defects              | Legacy TDM defects. Defects specific to PDH and SDH legacy error codes and defects.                                               | all             |

Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                                   | Level of Output |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| System defects          | Smart SFP device defects. Defects are specific to the Smart SFP you use.                                                                                                                                                                                            |                 |
| CoS queues              | Number of CoS queues configured.                                                                                                                                                                                                                                    | detail none     |
| Hold-Times              | Current interface hold-time up and hold-time down, in milliseconds.                                                                                                                                                                                                 | detail          |
| Current address         | Configured MAC address.                                                                                                                                                                                                                                             | detail none     |
| Hardware address        | Hardware MAC address.                                                                                                                                                                                                                                               | detail none     |
| Last flapped            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .         | detail none     |
| Statistics last cleared | Date, time, and how long ago the statistics for the interface were cleared. The format is <b>Statistics last cleared: <i>year-month-day hour.minute.second.timezone (hour.minute.second ago)</i></b> . For example, <b>2010-05-17 07:51:28 PDT (00:04:33 ago)</b> . | detail          |

Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Traffic statistics      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>                                    | detail          |
| IPv6 transit statistics | <p>If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface:</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |

Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)

| Field Name   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Input errors | <p data-bbox="431 369 1159 436">Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul data-bbox="431 474 1187 1808" style="list-style-type: none"> <li data-bbox="431 474 1110 541">• Errors—Sum of the incoming frame terminates and FCS errors.</li> <li data-bbox="431 579 1175 726">• Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li data-bbox="431 764 1162 831">• Framing errors—Number of packets received with an invalid frame checksum (FCS).</li> <li data-bbox="431 869 1167 936">• Runts—Number of frames received that are smaller than the runt threshold.</li> <li data-bbox="431 974 1182 1121">• Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li data-bbox="431 1159 1166 1348">• L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement.</li> <li data-bbox="431 1386 1183 1453">• L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li data-bbox="431 1491 1159 1600">• L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li data-bbox="431 1638 1170 1747">• FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li data-bbox="431 1785 938 1808">• Resource errors—Sum of transmit drops.</li> </ul> | detail none     |

Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Output errors | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• Errors—Sum of the outgoing frame terminates and FCS errors.</li> <li>• Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces.</li> </ul> | detail none     |

Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                       | Level of Output |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|                          | <ul style="list-style-type: none"> <li>• MTU errors—Number of packets whose size exceeded the MTU of the interface.</li> <li>• Resource errors—Sum of transmit drops.</li> </ul>                                                                                                                                        |                 |
| Egress queues            | Total number of egress queues supported on the specified interface.                                                                                                                                                                                                                                                     | detail          |
| Queue counters           | CoS queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> <li>• Queued packets—Number of queued packets.</li> <li>• Transmitted packets—Number of transmitted packets.</li> <li>• Dropped packets—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> | detail          |
| Input rate               | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                       | None specified  |
| Output rate              | Output rate in bps and pps.                                                                                                                                                                                                                                                                                             | None specified  |
| <b>Logical Interface</b> |                                                                                                                                                                                                                                                                                                                         |                 |
| Logical interface        | Name of the logical interface.                                                                                                                                                                                                                                                                                          | All levels      |
| Index                    | Index number of the logical interface, which reflects its initialization sequence.                                                                                                                                                                                                                                      | detail none     |
| SNMP ifIndex             | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                  | detail none     |

Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Flags              | <p>Information about the logical interface.</p> <p>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.</p>                                                                                      | All levels      |
| Input packets      | Number of packets received on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                   | detail none     |
| Output packets     | Number of packets transmitted on the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                | detail none     |
| Encapsulation      | <p>Encapsulation method used on the logical interface.</p> <ul style="list-style-type: none"> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul>                                                                                                                                                                                                                                            | All levels      |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |

Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)

| Field Name             | Field Description                                                                                                                                                                                                                                                                                                                                                                              | Level of Output |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Local statistics       | <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |
| Transit statistics     | <ul style="list-style-type: none"> <li>• Input bytes—Number of bytes received on the interface.</li> <li>• Output bytes—Number of bytes transmitted on the interface.</li> <li>• Input packets—Number of packets received on the interface.</li> <li>• Output packets—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p> | detail          |
| Addresses, Flags       | Information about the address flags.                                                                                                                                                                                                                                                                                                                                                           | detail none     |
| <i>protocol-family</i> | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.                                                                                                                                                                                                                                                       | brief           |
| MTU                    | Maximum transmission unit size on the physical interface.                                                                                                                                                                                                                                                                                                                                      | All levels      |
| Destination            | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                               | detail none     |
| Local                  | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                           | detail none     |

**Table 167: show interfaces smart-sfp-statistics Output Fields (Continued)**

| Field Name  | Field Description                                                                                                 | Level of Output |
|-------------|-------------------------------------------------------------------------------------------------------------------|-----------------|
| Broadcast   | Broadcast address of the logical interlace.                                                                       | detail none     |
| Generation  | Unique number for use by Juniper Networks technical support only.                                                 | detail          |
| Route table | Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0. | detail none     |

## Sample Output

### show interfaces smart-sfp-statistics (STM1)

```

user@host> show interfaces ge-5/3/9 smart-sfp-statistics
Physical interface: ge-5/3/9, Enabled, Physical link is Up
 Interface index: 300, SNMP ifIndex: 562
 Link-level type: Ethernet, MTU: 1518, MRU: 1526, 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
 Smart Transceiver Type: STM1/OC3
 Smart SFP: Configurable Type
 Smart SFP Configurations
 Loopback: None
 Mode: MEF8
 Smart SFP Ethernet port[P1] Statistics: Counters
 Rx frames 64376
 Rx bytes 4128352
 Rx errored fcs frames 0
 Rx length error frames 0
 Tx frames 2552954627

```

```

Tx bytes 120088140
Smart SFP AV IWF Encap/Decap Statistics: Counters
Rx Packets 0
Tx Packets 2552812339
Malformed Packets 0
Reordered Packets 0
Misordered Dropped Packets 0
Missing Packets 0
PlayedOut Packets 0
JitterBuffer Overrun Packets 0
JitterBuffer Underrun Packets 0
Smart SFP STM port[P0] statistics: Counters
Rx blerrors 0
Tx blerrors 0

Logical interface ge-5/3/9.32767 (Index 347) (SNMP ifIndex 659)
 Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [0x0000.0] Encapsulation: ENET2
 Input packets : 171346
 Output packets: 0
 Protocol multiservice, MTU: Unlimited
 Flags: Is-Primary

```

### show interfaces smart-sfp-statistics detail (STM1)

```

user@host> show interfaces ge-5/3/9 smart-sfp-statistics detail
Physical interface: ge-5/3/9, Enabled, Physical link is Up
 Interface index: 300, SNMP ifIndex: 562, Generation: 303
 Link-level type: Ethernet, MTU: 1518, MRU: 1526, 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
 Smart Transceiver Type: STM1/OC3
 Smart SFP: Configurable Type
 Smart SFP Configurations
 Loopback: None
 Mode: MEF8
 Smart SFP Ethernet port[P1] Statistics: Counters

```

```

Rx frames 105783
Rx bytes 6778400
Rx errored fcs frames 0
Rx length error frames 0
Tx frames 4209457628
Tx bytes 301542720
Smart SFP AV IWF Encap/Decap Statistics: Counters
Rx Packets 0
Tx Packets 4209223157
Malformed Packets 0
Reordered Packets 0
Misordered Dropped Packets 0
Missing Packets 0
PlayedOut Packets 0
JitterBuffer Overrun Packets 0
JitterBuffer Underrun Packets 0
Smart SFP STM port[P0] statistics: Counters
Rx blerrors 0
Tx blerrors 0

```

Logical interface ge-5/3/9.32767 (Index 347) (SNMP ifIndex 659) (Generation 186)

Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

Traffic statistics:

```

Input bytes : 23232960
Output bytes : 0
Input packets: 281769
Output packets: 0

```

Local statistics:

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Transit statistics:

```

Input bytes : 23232960 1008 bps
Output bytes : 0 0 bps
Input packets: 281769 1 pps
Output packets: 0 0 pps

```

Protocol multiservice, MTU: Unlimited, Generation: 227, Route table: 0

Flags: Is-Primary

Policer: Input: \_\_default\_arp\_policer\_\_

**show interfaces smart-sfp-statistics terse (STM1)**

```

user@host> show interfaces ge-5/3/9 smart-sfp-statistics terse
Interface Admin Link Proto Local Remote
ge-5/3/9 up up
ge-5/3/9.32767 up up multiservice

```

**show interfaces smart-sfp-statistics (DS3)**

```

user@host> show interfaces ge-5/3/0 smart-sfp-statistics
Physical interface: ge-5/3/0, Enabled, Physical link is Up
 Interface index: 291, SNMP ifIndex: 559
 Link-level type: Ethernet, MTU: 1518, MRU: 1526, 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
 Smart Transceiver Type: DS3
 Smart SFP: Configurable Type
 Smart SFP Configurations
 Loopback: None
 Encapsulation Circuit Id: 16
 Decapsulation Circuit Id: 16
 Mode: MEF8
 Dest MAC: 01:0f:25:00:00:03
 Smart SFP Ethernet port[P1] Statistics: Counters
 Rx frames 3013
 Rx bytes 220548
 Rx errored fcs frames 0
 Rx unicast frames 3013
 Rx multicast frames 0
 Rx broadcast frames 0
 Rx fragments 0
 Rx undersize frames 0
 Rx oversize frames 0
 Rx invalid vlan mismatch frames 0
 Tx frames 110842256

```

```

Tx bytes 4202709764
Tx unicast frames 18256
Tx multicast frames 110824000
Tx broadcast frames 0
Smart SFP AV IWF Encap/Decap Statistics: Counters
Rx Packets 0
Tx Packets 110824000
Malformed Packets 0
Reordered Packets 0
Misordered Dropped Packets 0
Missing Packets 0
PlayedOut Packets 0
JitterBuffer Overrun Packets 0
JitterBuffer Underrun Packets 0
Smart SFP DS3 port[P0] statistics: Counters
BiPolarVariations/Excessive zero errors 0
Tx B3 Errors 0
Code Violation path errors 0

Logical interface ge-5/3/0.32767 (Index 358) (SNMP ifIndex 660)
 Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [0x0000.0] Encapsulation: ENET2
 Input packets : 18341
 Output packets: 0
 Protocol multiservice, MTU: Unlimited
 Flags: None

```

### show interfaces smart-sfp-statistics detail (DS3)

```

Physical interface: ge-5/3/0, Enabled, Physical link is Up
 Interface index: 291, SNMP ifIndex: 559, Generation: 294
 Link-level type: Ethernet, MTU: 1518, MRU: 1526, 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
 Smart Transceiver Type: DS3
 Smart SFP: Configurable Type
 Smart SFP Configurations

```

Loopback: None  
 Encapsulation Circuit Id: 16  
 Decapsulation Circuit Id: 16  
 Mode: MEF8  
 Dest MAC: 01:0f:25:00:00:03

| Smart SFP Ethernet port[P1] Statistics: | Counters  |
|-----------------------------------------|-----------|
| Rx frames                               | 4329      |
| Rx bytes                                | 316992    |
| Rx errored fcs frames                   | 0         |
| Rx unicast frames                       | 4329      |
| Rx multicast frames                     | 0         |
| Rx broadcast frames                     | 0         |
| Rx fragments                            | 0         |
| Rx undersize frames                     | 0         |
| Rx oversize frames                      | 0         |
| Rx invalid vlan mismatch frames         | 0         |
| Tx frames                               | 159426251 |
| Tx bytes                                | 746264028 |
| Tx unicast frames                       | 26251     |
| Tx multicast frames                     | 159400000 |
| Tx broadcast frames                     | 0         |

| Smart SFP AV IWF Encap/Decap Statistics: | Counters  |
|------------------------------------------|-----------|
| Rx Packets                               | 0         |
| Tx Packets                               | 159400000 |
| Malformed Packets                        | 0         |
| Reordered Packets                        | 0         |
| Misordered Dropped Packets               | 0         |
| Missing Packets                          | 0         |
| PlayedOut Packets                        | 0         |
| JitterBuffer Overrun Packets             | 0         |
| JitterBuffer Underrun Packets            | 0         |

| Smart SFP DS3 port[P0] statistics:      | Counters |
|-----------------------------------------|----------|
| BiPolarVariations/Excessive zero errors | 0        |
| Tx B3 Errors                            | 0        |
| Code Violation path errors              | 0        |

Logical interface ge-5/3/0.32767 (Index 358) (SNMP ifIndex 660) (Generation 197)

Flags: Up SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

Traffic statistics:

|                |         |
|----------------|---------|
| Input bytes :  | 2985128 |
| Output bytes : | 0       |
| Input packets: | 26336   |

```

Output packets: 0
Local statistics:
 Input bytes : 0
Output bytes : 0
 Input packets: 0
 Output packets: 0
Transit statistics:
 Input bytes : 2985128 1024 bps
 Output bytes : 0 0 bps
 Input packets: 26336 1 pps
 Output packets: 0 0 pps
Protocol multiservice, MTU: Unlimited, Generation: 242, Route table: 0
Flags: None
Policer: Input: __default_arp_policer__

```

### show interfaces smart-sfp-statistics terse (DS3)

```

user@host> show interfaces ge-5/3/0 smart-sfp-statistics terse
Interface Admin Link Proto Local Remote
ge-5/3/0 up up
ge-5/3/0.32767 up up multiservice

```

### command-name

### Release Information

Command introduced in Junos OS Release 19.4.

### RELATED DOCUMENTATION

[Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | 246](#)

[ces-psn-channel \(tdm-options\) | 823](#)

[iwf-params \(tdm-options\) | 934](#)

[tdm-options \(Interfaces\) | 1176](#)

## show interfaces transport pm

### IN THIS SECTION

- [Syntax | 1737](#)
- [Description | 1737](#)
- [Options | 1737](#)
- [Required Privilege Level | 1738](#)
- [Output Fields | 1738](#)
- [Sample Output | 1740](#)
- [Release Information | 1747](#)

### Syntax

```
show interfaces transport pm (all | optics | otn) (all | current | currentday |
interval | previousday) (all | interface-name)
```

### Description

Display diagnostic data, warnings, and alarms for transport performance monitoring interfaces.

### Options

(all | optics | otn)—Display both optics and OTN information or either only optics or only OTN information.

(all | current | currentday | interval | previousday)—Display information for the current 15-minute interval, the current day, the ninety-six 15-minute intervals, and the previous day; information only for the current 15-minute interval; information only for the current 24 hours; information only for the ninety-six 15-minute intervals; information only for the previous day.

(all | *interface-name*)—Display information for all interfaces or only for the specified interface (for example, *et-fpc/pic/port*).

## Required Privilege Level

view

## Output Fields

Table 168 on page 1738 lists the output fields for the **show interfaces transport pm optics** command. Fields are listed in the approximate order in which they appear.

**Table 168: show interfaces transport pm Output Fields**

| Field Name                | Field Description                                                                                  |
|---------------------------|----------------------------------------------------------------------------------------------------|
| <b>Physical interface</b> | Name of the physical interface.                                                                    |
| <b>Interval</b>           | The 15 minute interval for performance monitoring.                                                 |
| <b>Suspect Flag</b>       | TRUE if the performance monitoring data for the inter                                              |
| <b>Reason</b>             | Reason for setting the suspect flag.                                                               |
| <b>COUNT</b>              | Measured value.                                                                                    |
| <b>THRESHOLD</b>          | Threshold value set.                                                                               |
| <b>TCA-ENABLED</b>        | Threshold crossing alert. Set to TRUE if enabled.                                                  |
| <b>TCA-RAISED</b>         | TRUE if enabled and the value crosses the threshold.                                               |
| <b>Near End PM</b>        | Near end threshold crossing defect trigger. For more in<br>see <a href="#">"tca" on page 757</a> . |
| <b>Far End PM</b>         | Far end threshold crossing defect trigger. For more info<br><a href="#">"tca" on page 757</a> .    |

Table 168: show interfaces transport pm Output Fields (Continued)

| Field Name                    | Field Description                                                                                 |
|-------------------------------|---------------------------------------------------------------------------------------------------|
| FEC PM                        | Forwarding equivalence class threshold crossing defect about each trigger, see "tca" on page 757. |
| BER PM                        | Bit error rate threshold crossing defect trigger. For more see "tca" on page 757.                 |
| CURRENT                       | Current value measured.                                                                           |
| PM                            | Performance monitor.                                                                              |
| MIN                           | Minimum value measured.                                                                           |
| MAX                           | Maximum value measured.                                                                           |
| AVG                           | Average value.                                                                                    |
| Lane Chromatic dispersion     | Residual chromatic dispersion measured.                                                           |
| Lane differential group delay | Measured differential group delay.                                                                |
|                               |                                                                                                   |
|                               |                                                                                                   |
| q Value                       | Measured Quality factor value.                                                                    |
| SNR                           | Signal to noise ratio.                                                                            |
| Tx output power               | The transmit laser output power.                                                                  |
| Rx input power                | The laser's received optical power.                                                               |

Table 168: show interfaces transport pm Output Fields (Continued)

| Field Name                     | Field Description                                                                                         |
|--------------------------------|-----------------------------------------------------------------------------------------------------------|
| Module temperature (Celsius)   | The laser's temperature.                                                                                  |
| Tx Laser bias current (0.1mA)  | Magnitude of the laser bias power setting current. The modulation of laser diodes and modulates currents. |
| Rx Laser bias current (0.1mA)  | Magnitude of the laser bias power setting current.                                                        |
| Carrier frequency offset (MHz) | Measured carrier frequency offset.                                                                        |

## Sample Output

### show interfaces transport pm

```

user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
 14:45-current Elapse time:900 Seconds
 Near End Suspect Flag:False Reason:None
 PM COUNT THRESHOLD TCA-ENABLED TCA-
 RAISED
 OTU-BBE 0 800 No No
 OTU-ES 0 135 No No
 OTU-SES 0 90 No No
 OTU-UAS 427 90 No No
 Far End Suspect Flag:True Reason:Unknown
 PM COUNT THRESHOLD TCA-ENABLED TCA-
 RAISED
 OTU-BBE 0 800 No No
 OTU-ES 0 135 No No
 OTU-SES 0 90 No No
 OTU-UAS 0 90 No No
 Near End Suspect Flag:False Reason:None
 PM COUNT THRESHOLD TCA-ENABLED TCA-
 RAISED
 ODU-BBE 0 800 No No

```

|                                                |             |                        |                |           |
|------------------------------------------------|-------------|------------------------|----------------|-----------|
| ODU-ES                                         | 0           | 135                    | No             | No        |
| ODU-SES                                        | 0           | 90                     | No             | No        |
| ODU-UAS                                        | 427         | 90                     | No             | No        |
| Far End                                        | Suspect     | Flag:True              | Reason:Unknown |           |
| PM                                             | COUNT       | THRESHOLD              | TCA-ENABLED    | TCA-      |
| RAISED                                         |             |                        |                |           |
| ODU-BBE                                        | 0           | 800                    | No             | No        |
| ODU-ES                                         | 0           | 135                    | No             | No        |
| ODU-SES                                        | 0           | 90                     | No             | No        |
| ODU-UAS                                        | 0           | 90                     | No             | No        |
| FEC                                            | Suspect     | Flag:False             | Reason:None    |           |
| PM                                             | COUNT       | THRESHOLD              | TCA-ENABLED    | TCA-      |
| RAISED                                         |             |                        |                |           |
| FEC-CorrectedErr                               | 2008544300  | 0                      | NA             | NA        |
| FEC-UncorrectedWords                           | 0           | 0                      | NA             | NA        |
| BER                                            | Suspect     | Flag:False             | Reason:None    |           |
| PM                                             | MIN         | MAX                    | AVG            | THRESHOLD |
| ENABLED TCA-RAISED                             |             |                        |                |           |
| BER                                            | 3.6e-5      | 5.8e-5                 | 3.6e-5         | 10.0e-3   |
| No                                             | Yes         |                        |                |           |
| Physical interface: et-0/1/0, SNMP ifIndex 515 |             |                        |                |           |
| 14:45-current                                  |             |                        |                |           |
| Suspect                                        | Flag:True   | Reason:Object Disabled |                |           |
| PM                                             | CURRENT     | MIN                    | MAX            | AVG       |
| THRESHOLD                                      | TCA-ENABLED | TCA-RAISED             |                |           |
| (MAX)                                          | (MIN)       | (MAX)                  | (MIN)          | (MAX)     |
| Lane chromatic dispersion                      | 0           | 0                      | 0              | 0         |
| 0                                              | NA          | NA                     | NA             | NA        |
| Lane differential group delay                  | 0           | 0                      | 0              | 0         |
| 0                                              | NA          | NA                     | NA             | NA        |
| q Value                                        | 120         | 120                    | 120            | 120       |
| 0                                              | NA          | NA                     | NA             | NA        |
| SNR                                            | 28          | 28                     | 29             | 28        |
| 0                                              | NA          | NA                     | NA             | NA        |
| Tx output power(0.01dBm)                       | -5000       | -5000                  | -5000          | -5000     |
| -100                                           | No          | No                     | No             | No        |
| Rx output power(0.01dBm)                       | -3642       | -3665                  | -3626          | -3637     |
| -500                                           | No          | No                     | No             | No        |
| Module temperature(Celsius)                    | 46          | 46                     | 46             | 46        |
| 75                                             | No          | No                     | No             | No        |
| Tx laser bias current(0.1mA)                   | 0           | 0                      | 0              | 0         |
| 0                                              | NA          | NA                     | NA             | NA        |

|                               |      |      |      |      |       |
|-------------------------------|------|------|------|------|-------|
| Rx laser bias current(0.1mA)  | 1270 | 1270 | 1270 | 1270 | 0     |
| 0                             | NA   | NA   | NA   | NA   |       |
| Carrier frequency offset(MHz) | -186 | -186 | -186 | -186 | -5000 |
| 5000                          | No   | No   | No   | No   |       |

### show interfaces transport (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```

user@host > show interfaces transport et-3/0/0
Administrative State: In Service
Operational State: Normal

```

### show interfaces transport pm (MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

```

user@host > show interfaces transport pm otn current et-3/0/0
Physical interface: et-3/0/0, SNMP ifIndex 564
23:30-current Elapsed time:455 Seconds
Near End Suspect Flag:False Reason:Not Applicable
 PM COUNT THRESHOLD TCA-ENABLED TCA-
RAISED
 OTU-BBE 0 800 No No
 OTU-ES 0 135 No No
 OTU-SES 0 90 No No
 OTU-UAS 0 90 No No
Far End Suspect Flag:False Reason:Not Applicable
 PM COUNT THRESHOLD TCA-ENABLED TCA-
RAISED
 OTU-BBE 0 800 No No
 OTU-ES 0 135 No No
 OTU-SES 0 90 No No
 OTU-UAS 0 90 No No
Near End Suspect Flag:False Reason:Not Applicable
 PM COUNT THRESHOLD TCA-ENABLED TCA-
RAISED
 ODU-BBE 0 800 No No
 ODU-ES 0 135 No No
 ODU-SES 0 90 No No
 ODU-UAS 0 90 No No
Far End Suspect Flag:False Reason:Not Applicable

```

| PM                   | COUNT              | THRESHOLD | TCA-ENABLED           | TCA-      |      |
|----------------------|--------------------|-----------|-----------------------|-----------|------|
| RAISED               |                    |           |                       |           |      |
| ODU-BBE              | 0                  | 800       | No                    | No        |      |
| ODU-ES               | 0                  | 135       | No                    | No        |      |
| ODU-SES              | 0                  | 90        | No                    | No        |      |
| ODU-UAS              | 0                  | 90        | No                    | No        |      |
| FEC                  | Suspect Flag:False |           | Reason:Not Applicable |           |      |
| PM                   | COUNT              | THRESHOLD | TCA-ENABLED           | TCA-      |      |
| RAISED               |                    |           |                       |           |      |
| FEC-CorrectedErr     | 30865849           | 0         | NA                    | NA        |      |
| FEC-UncorrectedWords | 0                  | 0         | NA                    | NA        |      |
| BER                  | Suspect Flag:False |           | Reason:Not Applicable |           |      |
| PM                   | MIN                | MAX       | AVG                   | THRESHOLD | TCA- |
| ENABLED TCA-RAISED   |                    |           |                       |           |      |
| BER                  | 4.0e-7             | 5.9e-7    | 5.1e-7                | 1.0e-2    |      |
| No                   | No                 |           |                       |           |      |

## command-name

```

user@host > show interfaces transport pm optics current et-2/0/0
Physical interface: et-3/0/0, SNMP ifIndex 564
23:30-current
Suspect Flag:True Reason:Not Applicable

```

| PM                                | CURRENT     | MIN        | MAX   | AVG   |       |
|-----------------------------------|-------------|------------|-------|-------|-------|
| THRESHOLD                         | TCA-ENABLED | TCA-RAISED |       |       | (MIN) |
| (MAX)                             | (MIN)       | (MAX)      | (MIN) | (MAX) |       |
| Lane chromatic dispersion(ps/nm)  | 0           | 0          | 0     | 51    | 0     |
| 0                                 | NA          | NA         | NA    | NA    |       |
| Lane differential group delay(ps) | -13         | 13         | 0     | 11    | 0     |
| 0                                 | NA          | NA         | NA    | NA    |       |
| q Value(0.1dB)                    | 0           | -1         | 5     | 137   | 0     |
| 0                                 | NA          | NA         | NA    | NA    |       |
| SNR(0.1dB)                        | 137         | 138        | 137   | 86    | 0     |
| 0                                 | NA          | NA         | NA    | NA    |       |
| Tx output power(0.01dBm)          | 83          | 95         | 83    | 142   | -300  |
| -100                              | No          | No         | No    | No    |       |
| Rx input power(0.01dBm)           | 141         | 142        | 141   | 106   | -1800 |
| -500                              | No          | No         | No    | No    |       |
| Module temperature(Celsius)       | 106         | 109        | 106   | -31   | -5    |
| 75                                | No          | No         | No    | No    |       |

```

Tx laser bias current(0.1mA) -31 0 0 38 0
0 NA NA NA NA
Rx laser bias current(0.1mA) 38 38 38 0 0
0 NA NA NA NA
Carrier frequency offset(MHz) 0 0 2 0 -5000
5000 No No No No

```

### show interfaces transport (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```

user@host > show interfaces transport et-8/0/0
Administrative State: In Service
Operational State: Normal

```

### show interfaces transport pm optics (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```

user@host > show interfaces transport pm optics current et-4/0/0
Physical interface: et-4/0/0, SNMP ifIndex 544
02:45-current
Suspect Flag:False Reason:Not Applicable
 PM CURRENT MIN MAX AVG
THRESHOLD TCA-ENABLED TCA-RAISED
(MIN) (MIN) (MAX) (MIN) (MAX)
Lane chromatic dispersion(ps/nm) -6 -32 45 -1 0
0 NA NA NA NA
Lane differential group delay(ps) 3 2 4 3 0
0 NA NA NA NA
Lane Q2 factor(0.1dB) 154 154 155 154 0
0 NA NA NA NA
SNR(0.1dB) 167 164 171 165 0
0 NA NA NA NA
Carrier frequency offset(MHz) 0 0 0 0 -3600
3600 No No No No
Tx output power(0.01dBm) 0 0 0 0 -1100
300 No No No No
Rx input total power(0.01dBm) 0 0 0 0 -3000
300 No No No No

```

```

Module temperature(Celsius) 53 53 55 53 -5
75 No No No No

```

### show interfaces transport pm otn (PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC)

```

user@host> show interfaces transport pm otn previousday et-4/0/0
Physical interface: et-4/0/0, SNMP ifIndex 544
02:45-current
Suspect Flag:False Reason:Not Applicable
PM CURRENT MIN MAX AVG
THRESHOLD TCA-ENABLED TCA-RAISED
(MIN) (MIN) (MAX)
(MAX) (MIN) (MAX) (MIN) (MAX)
Lane chromatic dispersion(ps/nm) -6 -32 45 -1 0
0 NA NA NA NA
Lane differential group delay(ps) 3 2 4 3 0
0 NA NA NA NA
Lane Q2 factor(0.1dB) 154 154 155 154 0
0 NA NA NA NA
SNR(0.1dB) 167 164 171 165 0
0 NA NA NA NA
Carrier frequency offset(MHz) 0 0 0 0 -3600
3600 No No No No
Tx output power(0.01dBm) 0 0 0 0 -1100
300 No No No No
Rx input total power(0.01dBm) 0 0 0 0 -3000
300 No No No No
Module temperature(Celsius) 53 53 55 53 -5
75 No No No No

```

### show interfaces transport pm optics (MX10003, MX10008, MX10016, and MX204 router with QSFP-100GE-DWDM2 transceiver)

```

user@host> show interfaces transport pm optics current et-2/0/0

Physical interface: et-2/0/0, SNMP ifIndex 934
18:00-current
Suspect Flag:False Reason:Not Applicable

```

| PM THRESHOLD                  | TCA-ENABLED |       | CURRENT TCA-RAISED |         | MIN     | MAX     | AVG     |       |
|-------------------------------|-------------|-------|--------------------|---------|---------|---------|---------|-------|
| (MAX)                         | (MIN)       | (MAX) | (MIN)              | (MAX)   |         |         | (MIN)   |       |
| Module temperature (Celsius)  |             |       |                    | 32      | 32      | 32      | 32      | -5    |
| 75                            | No          | No    | No                 | No      |         |         |         |       |
| Pre-FEC BER                   |             |       |                    | 0.00e-0 | 0.00e-0 | 0.00e-0 | 0.00e-0 | 0     |
| 1.00e-3                       | NA          | No    | NA                 | No      |         |         |         |       |
| Uncorrected FER               |             |       |                    | 0.00e-0 | 0.00e-0 | 0.00e-0 | 0.00e-0 | 0     |
| 0.00e-0                       | NA          | No    | NA                 | No      |         |         |         |       |
| Lane 0                        |             |       |                    |         |         |         |         |       |
| PM THRESHOLD                  | TCA-ENABLED |       | CURRENT TCA-RAISED |         | MIN     | MAX     | AVG     |       |
| (MAX)                         | (MIN)       | (MAX) | (MIN)              | (MAX)   |         |         | (MIN)   |       |
| SNR (0.1dB)                   |             |       |                    | 0       | 0       | 0       | 0       | 150   |
| 0                             | No          | NA    | No                 | NA      |         |         |         |       |
| Laser tx power (0.01dBm)      |             |       |                    | -951    | -966    | -943    | -943    | -300  |
| -100                          | No          | No    | No                 | No      |         |         |         |       |
| Laser rx power (0.01dBm)      |             |       |                    | 8       | 6       | 16      | 6       | -1800 |
| -500                          | No          | No    | No                 | No      |         |         |         |       |
| Tx laser bias current (0.1mA) |             |       |                    | 691     | 675     | 713     | 675     | 0     |
| 0                             | NA          | NA    | NA                 | NA      |         |         |         |       |
| Laser frequency Error (MHz)   |             |       |                    | -389    | -429    | -339    | -353    | -3000 |
| 3000                          | No          | No    | No                 | No      |         |         |         |       |
| TEC Current (0.1mA)           |             |       |                    | 492     | 423     | 509     | 435     | -5000 |
| 5000                          | No          | No    | No                 | No      |         |         |         |       |
| Residual ISI (0.1ps/nm)       |             |       |                    | 0       | 0       | 0       | 0       | -800  |
| 800                           | No          | No    | No                 | No      |         |         |         |       |
| PAM Histogram                 |             |       |                    | 0       | 0       | 0       | 0       | 0     |
| 16384                         | NA          | No    | NA                 | No      |         |         |         |       |
| Lane 1                        |             |       |                    |         |         |         |         |       |
| PM THRESHOLD                  | TCA-ENABLED |       | CURRENT TCA-RAISED |         | MIN     | MAX     | AVG     |       |
| (MAX)                         | (MIN)       | (MAX) | (MIN)              | (MAX)   |         |         | (MIN)   |       |
| SNR (0.1dB)                   |             |       |                    | 0       | 0       | 111     | 0       | 150   |
| 0                             | No          | NA    | No                 | NA      |         |         |         |       |
| Laser tx power (0.01dBm)      |             |       |                    | -954    | -956    | -943    | -943    | -300  |
| -100                          | No          | No    | No                 | No      |         |         |         |       |
| Laser rx power (0.01dBm)      |             |       |                    | 27      | 14      | 28      | 14      | -1800 |
| -500                          | No          | No    | No                 | No      |         |         |         |       |
| Tx laser bias current (0.1mA) |             |       |                    | 781     | 770     | 797     | 770     | 0     |

|       |                             |    |    |    |     |     |     |     |       |
|-------|-----------------------------|----|----|----|-----|-----|-----|-----|-------|
| 0     | NA                          | NA | NA | NA |     |     |     |     |       |
|       | Laser frequency Error (MHz) |    |    |    | 0   | -59 | 19  | 0   | -3000 |
| 3000  | No                          | No | No | No |     |     |     |     |       |
|       | TEC Current (0.1mA)         |    |    |    | 486 | 452 | 534 | 475 | -5000 |
| 5000  | No                          | No | No | No |     |     |     |     |       |
|       | Residual ISI (0.1ps/nm)     |    |    |    | 0   | 0   | 0   | 0   | -800  |
| 800   | No                          | No | No | No |     |     |     |     |       |
|       | PAM Histogram               |    |    |    | 0   | 0   | 0   | 0   | 0     |
| 16384 | NA                          | No | NA | No |     |     |     |     |       |

## Release Information

Command introduced in Junos OS Release 14.2.

## RELATED DOCUMENTATION

[clear interfaces transport pm | 1236](#)

[tca | 757](#)

[transport-monitoring | 761](#)

## show l2-learning instance

### IN THIS SECTION

- [Syntax | 1748](#)
- [Description | 1748](#)
- [Options | 1748](#)
- [Required Privilege Level | 1748](#)
- [Output Fields | 1748](#)
- [Sample Output | 1749](#)
- [Release Information | 1749](#)

## Syntax

```
show l2-learning instance
```

## Description

Display Layer 2 learning properties for all the configured routing instances.

## Options

This command has no options.

## Required Privilege Level

view

## Output Fields

[Table 169 on page 1748](#) describes the output fields for the **show l2-learning instance** command. Output fields are listed in the approximate order in which they appear.

**Table 169: show l2-learning instance Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                   |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Routing Instance</b> | Name of routing instance.                                                                                                                                                                                                                           |
| <b>Bridging Domain</b>  | Name of bridging domain.<br><br>On MX Series routers you can use the <b>show l2-learning instance &lt;extensive&gt;</b> command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID. |
| <b>Index</b>            | Number associated with the routing instance or bridging domain.                                                                                                                                                                                     |
| <b>Logical System</b>   | Name of logical system or <b>Default</b> if no logical system is configured.                                                                                                                                                                        |

Table 169: show l2-learning instance Output Fields (Continued)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Routing instance flags</b> | <p>Status of Layer 2 learning properties for each routing instance:</p> <ul style="list-style-type: none"> <li>• <b>DL</b>—MAC learning is disabled.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>AD</b>—Packets are dropped after MAC address limit is reached.</li> <li>• <b>LH</b>—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses.</li> </ul> |
| <b>MAC limit</b>              | Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain.                                                                                                                                                                                                                                                                                                                                                     |

## Sample Output

### show l2-learning instance

```

user@host> show l2-learning instance
Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
 AD -packet action drop, LH -mac limit hit)

Routing Bridging Index Logical Routing MAC
Instance Domain System flags limit
__juniper_private1__
vs1 vlan100 3 Default 5120
vs1 vlan200 4 Default 5120

```

## Release Information

(MX Series routers only) Command introduced in Junos OS Release 8.4.

## show l2-learning redundancy-groups

### IN THIS SECTION

- [Syntax | 1750](#)
- [Description | 1750](#)
- [Options | 1750](#)
- [Required Privilege Level | 1751](#)
- [Output Fields | 1751](#)
- [Sample Output | 1753](#)
- [Release Information | 1756](#)

### Syntax

```
show l2-learning redundancy-groups
logical-system [system-name | all]
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs
```

### Description

(MX Series routers only) Display ARP statistics, Neighbor Discovery statistics, or remote MAC addresses for the Multi-Chassis Aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a router or switch or logical systems on a router or switch. Note that the Redundancy Group ID is inherited by the bridging domain or VLAN from member AE interfaces.

### Options

**logical-system** [*system-name* | all] (Optional) Display information for a specified logical system or all systems.

|                            |                                                                                                                                                                                                              |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>redundancy-group-id</i> | (Optional) The redundancy group identification number. The Inter-Chassis Control Protocol (ICCP) uses the redundancy group ID to associate the routing or switching devices contained in a redundancy group. |
| <i>arp-statistics</i>      | (Optional) Count of ARP packets sent and received by the two MC-AE nodes.                                                                                                                                    |
| <i>nd-statistics</i>       | (Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.                                                                                                                     |
| <i>remote-macs</i>         | (Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.                                                                                                     |

## Required Privilege Level

view

## Output Fields

Output fields are listed in the approximate order in which they appear.

**Table 170: show l2-learning redundancy-groups arp-statistics Output Fields**

| Field Name                    | Field Description                                                              |
|-------------------------------|--------------------------------------------------------------------------------|
| Redundancy Group ID           | Redundancy Group to which the following details apply.                         |
| MCLAG ARP Statistics Group ID | ARP statistics for this Multichassis Link Aggregation Group (MC-LAG) instance. |
| ARP Rx Count From Line        | Total number of ARPs received from the Line.                                   |
| ARP Tx Count To Peer          | Total number of ARPs sent to the peer.                                         |
| ARP Rx Count From Peer        | Total number of ARPs received from the peer.                                   |

**Table 170: show I2-learning redundancy-groups arp-statistics Output Fields (Continued)**

| Field Name                        | Field Description                                         |
|-----------------------------------|-----------------------------------------------------------|
| ARP Drop Count received from line | Total number of ARPs sent by the peer that were received. |
| ARP Drop Count received from peer | Total number of ARPs sent by the peer that were dropped   |
| Service-id                        | Service ID (configured at the routing instance level).    |

**Table 171: show I2-learning redundancy-groups nd-statistics Output Fields**

| Field Name                       | Field Description                                                                            |
|----------------------------------|----------------------------------------------------------------------------------------------|
| Redundancy Group ID              | Redundancy Group to which the following details apply.                                       |
| MCLAG ND Statistics Group ID     | Neighbor Discovery statistics for this Multichassis Link Aggregation Group (MCLAG) instance. |
| ND Rx Count From Line            | Total number of Neighbor Discovery packets received from the Line.                           |
| ND Tx Count To Peer              | Total number of Neighbor Discovery packets sent to the peer.                                 |
| ND Rx Count From Peer            | Total number of Neighbor Discovery packets received from the peer.                           |
| ND Drop Count received from line | Total number of Neighbor Discovery packets sent by the peer that were received.              |

**Table 171: show I2-learning redundancy-groups nd-statistics Output Fields (Continued)**

| Field Name                       | Field Description                                                             |
|----------------------------------|-------------------------------------------------------------------------------|
| ND Drop Count received from peer | Total number of Neighbor Discovery packets sent by the peer that were dropped |
| Service-id                       | Service ID (configured at the routing instance level).                        |

**Table 172: show I2-learning redundancy-groups remote-macs Output Fields**

| Field Name          | Field Description                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------------|
| Redundancy Group ID | Redundancy Group to which the following details apply.                                                             |
| Peer-Addr           | IP address of the remote peer.                                                                                     |
| VLAN                | Virtual LAN identifier associated with the redundancy group.                                                       |
| MAC                 | Hardware media access control address associated with the redundancy group.                                        |
| MCAE-ID             | ID number of the MC-AE used by the redundancy group.                                                               |
| Flags               | Connection state: local connect or Remote connect. If no flag is shown, the redundancy group may not be connected. |
| Status              | Installation state: Installed or Not Installed.                                                                    |

## Sample Output

### show I2-learning redundancy-groups arp-statistics

```
user@host> show I2-learning redundancy-groups arp-statistics
Logical System : default
```

```
Redundancy Group ID : 1 Flags : Local Connect, Remote Connect
```

#### MCLAG ARP Statistics

```
Group ID : 1
ARP Rx Count From Line : 52
ARP Tx Count To Peer : 15
ARP Rx Count From Peer : 39
ARP Install Count : 34
ARP Drop Count received from line : 37
ARP Drop Count received from peer : 5
```

### show l2-learning redundancy-groups nd-statistics

```
user@host> show l2-learning redundancy-groups nd-statistics
```

```
Logical System : default
```

```
Redundancy Group ID : 1 Flags : Local Connect, Remote Connect
```

#### MCLAG ND Statistics

```
Group ID : 1
ND Rx Count From Line : 52
ND Tx Count To Peer : 15
ND Rx Count From Peer : 39
ND Install Count : 34
ND Drop Count received from line : 37
ND Drop Count received from peer : 5
```

### show l2-learning redundancy-groups remote-macs

```
user@host> show l2-learning redundancy-groups <redundancy-group-id> remote-macs
```

```
Redundancy Group ID : 1 Flags : Local Connect, Remote Connect
```

| Service-id | Peer-Addr | VLAN | MAC               | MCAE-ID | Subunit | Opcode |
|------------|-----------|------|-------------------|---------|---------|--------|
| Flags      | Status    |      |                   |         |         |        |
| 10         | 10.1.1.2  | 100  | 64:87:88:6a:df:f0 | 1       | 0       | 1      |
| 0          | Installed |      |                   |         |         |        |

**show l2-learning redundancy-groups logical-system arp-statistics (for Logical Systems)**

```
user@host> show l2-learning redundancy-groups logical-system LS1 arp-statistics

Redundancy Group ID : 1 Flags : Local Connect, Remote Connect

MCLAG ARP Statistics
Group ID : 1
ARP Rx Count From Line : 52
ARP Tx Count To Peer : 15
ARP Rx Count From Peer : 39
ARP Install Count : 34
ARP Drop Count received from line : 37
ARP Drop Count received from peer : 5
```

**show l2-learning redundancy-groups logical-system nd-statistics (for Logical Systems)**

```
user@host> show l2-learning redundancy-groups logical-system LS1 nd-statistics

Redundancy Group ID : 1 Flags : Local Connect, Remote Connect

MCLAG ND Statistics
Group ID : 1
ND Rx Count From Line : 52
ND Tx Count To Peer : 15
ND Rx Count From Peer : 39
ND Install Count : 34
ND Drop Count received from line : 37
ND Drop Count received from peer : 5
```

**show l2-learning redundancy-groups group-id**

```
user@host> show l2-learning redundancy-groups 1

Redundancy Group ID : 1 Flags : Local Connect, Remote Connect
```

## show l2-learning redundancy-groups logical-system

```
user@host> show l2-learning redundancy-groups logical-system ls1

Redundancy Group ID : 2 Flags : Local Connect,Remote Connect
```

### Release Information

Command introduced in Junos OS Release 13.2.

Support for logical systems added in Junos OS Release 14.1.

### RELATED DOCUMENTATION

[Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up](#)

[show interfaces mc-ae | 1688](#)

[Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers and QFX Series Switches](#)

[Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up](#)

## show lacp interfaces

### IN THIS SECTION

- [Syntax | 1757](#)
- [Description | 1757](#)
- [Options | 1757](#)
- [Required Privilege Level | 1757](#)
- [Output Fields | 1757](#)
- [Sample Output | 1763](#)
- [Release Information | 1764](#)

## Syntax

```
show lacp interfaces
<interface-name>
extensive
```

## Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interface.

## Options

|                              |                                                                                                                                                                                                                                                                                                                                                |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>none</b>                  | Display LACP information for all interfaces.                                                                                                                                                                                                                                                                                                   |
| <b><i>interface-name</i></b> | (Optional) Display LACP information for the specified interface: <ul style="list-style-type: none"> <li>• Aggregated Ethernet—<b><i>ae-number</i></b></li> <li>• Fast Ethernet—<b><i>fe-fpc/pic/port</i></b></li> <li>• Gigabit Ethernet—<b><i>ge-fpc/pic/port</i></b></li> <li>• 10 Gigabit Ethernet—<b><i>xe-fpc/pic/port</i></b></li> </ul> |
| <b>extensive</b>             | Display LACP information for the interface in detail.                                                                                                                                                                                                                                                                                          |

**NOTE:** The **show lacp interfaces** command returns the following error message if your system is not configured in either active or passive LACP mode:  
“Warning: lacp subsystem not running – not needed by configuration”

## Required Privilege Level

view

## Output Fields

[Table 173 on page 1758](#) lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 173: show lacp interfaces Output Fields**

| Field Name                  | Field Description                                                                                                    | Level of Output |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------|-----------------|
| LACP State                  | For a child interface configured with the force-up statement, LACP state displays FUP along with the interface name. | All Levels      |
| <b>Aggregated interface</b> | Aggregated interface value.                                                                                          | All Levels      |

Table 173: show lacp interfaces Output Fields (Continued)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>LACP State</b> | <p>LACP state information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>• <b>Role</b>—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> <li>• <b>Actor</b>—Local device participating in LACP negotiation.</li> <li>• <b>Partner</b>—Remote device participating in LACP negotiation.</li> </ul> </li> <li>• <b>Exp</b>—Expired state. <b>Yes</b> indicates the actor or partner is in an expired state. <b>No</b> indicates the actor or partner is not in an expired state.</li> <li>• <b>Def</b>—Default. <b>Yes</b> indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. <b>No</b> indicates the operational partner information in use has been received in an LACP PDU.</li> <li>• <b>Dist</b>—Distribution of outgoing frames. <b>No</b> indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is <b>Yes</b>.</li> <li>• <b>Col</b>—Collection of incoming frames. <b>Yes</b> indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is <b>No</b>.</li> <li>• <b>Syn</b>—Synchronization. If the value is <b>Yes</b>, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is <b>No</b>, the link is not synchronized. It is currently not in the right aggregation.</li> <li>• <b>Aggr</b>—Ability of aggregation port to aggregate (<b>Yes</b>) or to operate only as an individual link (<b>No</b>).</li> <li>• <b>Timeout</b>—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (<b>Slow Timeout</b> or <b>Fast Timeout</b>). In a fast timeout, PDUs are sent every second and in a slow timeout, PDUs are sent every 30 seconds. LACP timeout occurs when 3 consecutive PDUs are missed. If LACP</li> </ul> | All Levels      |

Table 173: show lacp interfaces Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <p>timeout is a fast timeout, the time taken when 3 consecutive PDUs are missed is 3 seconds (3x1 second). If LACP timeout is a slow timeout, the time taken is 90 seconds( 3x30 seconds).</p> <ul style="list-style-type: none"> <li>• <b>Activity</b>—Actor or partner's port activity. <b>Passive</b> indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is <b>Active</b>. <b>Active</b> indicates the port's preference to participate in the protocol regardless of the partner's control value.</li> <li>• Core isolation state down (CDN)— LACP interface state. Down indicates the LACP interface is down because all the eBGP sessions for Ethernet VPN (EVPN) are down.</li> </ul> |                 |

Table 173: show lacp interfaces Output Fields (Continued)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| LACP Protocol | <p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.</li> <li>• <b>Receive State</b>—One of the following values: <ul style="list-style-type: none"> <li>• <b>Current</b>—The state machine receives an LACP PDU and enters the <b>Current</b> state.</li> <li>• <b>Defaulted</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires a second time, the state machine enters the <b>Defaulted</b> state.</li> <li>• <b>Expired</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires once, the state machine enters the <b>Expired</b> state.</li> <li>• <b>Initialize</b>—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the <b>Initialize</b> state.</li> <li>• <b>LACP Disabled</b>—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to <b>LACP Disabled</b>. This state is similar to the <b>Defaulted</b> state, except that the port is forced to operate as an individual port.</li> <li>• <b>Port Disabled</b>—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the <b>Port Disabled</b> state.</li> </ul> </li> <li>• <b>Transmit State</b>—Transmit state of state machine. One of the following values: <ul style="list-style-type: none"> <li>• <b>Fast Periodic</b>—Periodic transmissions are enabled at a fast transmission rate.</li> <li>• <b>No Periodic</b>—Periodic transmissions are disabled.</li> <li>• <b>Periodic Timer</b>—Transitory state entered when the periodic timer expires.</li> <li>• <b>Slow Periodic</b>—Periodic transmissions are enabled at a slow transmission rate.</li> </ul> </li> <li>• <b>Mux State</b>—State of the multiplexer state machine for the aggregation port. The state is one of the following values:</li> </ul> | All Levels      |

Table 173: show lacp interfaces Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <li>• <b>Attached</b>—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.</li> <li>• <b>Collecting</b>—<b>Yes</b> indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. <b>No</b> indicates the receive function of this link is not enabled.</li> <li>• <b>Collecting Distributing</b>—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.</li> <li>• <b>Detached</b>—Process of detaching the port from the aggregator is in progress.</li> <li>• <b>Distributing</b>—<b>Yes</b> indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. <b>No</b> indicates the transmit function of this link is not enabled.</li> <li>• <b>Waiting</b>—Multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul> |                 |

Table 173: show lacp interfaces Output Fields (Continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| LACP info  | <ul style="list-style-type: none"> <li>• <b>Role</b> can be one of the following: <ul style="list-style-type: none"> <li>• <b>Actor</b>—Local device participating in LACP negotiation.</li> <li>• <b>Partner</b>—Remote device participating in LACP negotiation.</li> </ul> </li> <li>• <b>System priority</b>—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>System identifier</b>—Actor or partner system ID, encoded as a MAC address.</li> <li>• <b>Port priority</b>—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.</li> <li>• <b>Port number</b>—Port number assigned to the port by the actor or partner, encoded as an unsigned integer.</li> <li>• <b>Port key</b>—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.</li> </ul> | Extensive       |

## Sample Output

### show lacp interfaces (Aggregated Ethernet)

```

user@host> show lacp interfaces ae0 extensive
LACP state: Role Exp Def Dist Col Syn Aggr Timeout Activity
ge-0/0/1 Actor No No Yes Yes Yes Yes Fast Active
ge-0/0/1 Partner No No Yes Yes Yes Yes Fast Active
ge-0/0/2 Actor No No Yes Yes Yes Yes Fast Active
ge-0/0/2 Partner No No Yes Yes Yes Yes Fast Active
ge-0/0/3 Actor No No Yes Yes Yes Yes Fast Active
ge-0/0/3 Partner No No Yes Yes Yes Yes Fast Active
LACP protocol: Receive State Transmit State Mux State
ge-0/0/1 Current Fast periodic Collecting distributing
ge-0/0/2 Current Fast periodic Collecting distributing
ge-0/0/3 Current Fast periodic Collecting distributing
LACP info: Role System System Port Port
Port

```

| key |          |         | priority | identifier        | priority | number |
|-----|----------|---------|----------|-------------------|----------|--------|
| 1   | ge-0/0/1 | Actor   | 127      | 00:05:86:4e:b6:c0 | 127      |        |
| 1   | 1        |         |          |                   |          |        |
| 1   | ge-0/0/1 | Partner | 127      | 00:05:86:7e:d3:c0 | 127      |        |
| 1   | 1        |         |          |                   |          |        |
| 2   | ge-0/0/2 | Actor   | 127      | 00:05:86:4e:b6:c0 | 127      |        |
| 2   | 1        |         |          |                   |          |        |
| 2   | ge-0/0/2 | Partner | 127      | 00:05:86:7e:d3:c0 | 127      |        |
| 2   | 1        |         |          |                   |          |        |
| 3   | ge-0/0/3 | Actor   | 127      | 00:05:86:4e:b6:c0 | 127      |        |
| 3   | 1        |         |          |                   |          |        |
| 3   | ge-0/0/3 | Partner | 127      | 00:05:86:7e:d3:c0 | 127      |        |
| 3   | 1        |         |          |                   |          |        |

### show lacp interfaces (Gigabit Ethernet)

```

user@host> show lacp interfaces ge-0/3/0
Aggregated interface: ae0
 LACP State: Role Exp Def Dist Col Syn Aggr Timeout Activity
 ge-0/3/0 Actor No No Yes Yes Yes Yes Fast Active
 ge-0/3/0 Partner No No Yes Yes Yes Yes Fast Active
 LACP Protocol: Receive State Transmit State Mux State
 ge-0/3/0 Current Fast periodic Collecting distributing

```

### show lacp interfaces (10 Gigabit Ethernet)

```

user@host> show lacp interfaces xe-1/0/2
Aggregated interface: ae0
 LACP State: Role Exp Def Dist Col Syn Aggr Timeout Activity
 xe-1/0/2 Actor No No Yes Yes Yes Yes Fast Active
 xe-1/0/2 Partner No No Yes Yes Yes Yes Fast Active
 LACP Protocol: Receive State Transmit State Mux State
 xe-1/0/2 Current Fast periodic Collecting distributing

```

## Release Information

Command introduced in Junos OS Release 7.6.

**extensive** statement introduced in Junos OS Release 16.1R1.

## RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)

*Configuring Link Aggregation*

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

*Understanding Aggregated Ethernet Interfaces and LACP for Switches*

[Junos OS Interfaces Fundamentals Configuration Guide](#)

## show pppoe interfaces

### IN THIS SECTION

- [Syntax | 1765](#)
- [Description | 1766](#)
- [Options | 1766](#)
- [Required Privilege Level | 1766](#)
- [Output Fields | 1766](#)
- [Sample Output | 1769](#)
- [Release Information | 1771](#)

## Syntax

```
show pppoe interfaces
<brief | detail
<pp0.logical>
```

## Description

Display session-specific information about PPPoE interfaces.

## Options

**none** Display interface information for all PPPoE interfaces.

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

**pp0.logical** (Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16385. The logical unit number for dynamic interfaces can be a value from 1073741824 through the maximum number of logical interfaces supported on your router.

## Required Privilege Level

view

## Output Fields

Table 174 on page 1766 lists the output fields for the **show pppoe interfaces** command. Output fields are listed in the approximate order in which they appear. Not all fields are displayed for PPPoE interfaces on M120 and M320 routers in server mode.

**Table 174: show pppoe interfaces Output Fields**

| Field Name               | Field Description                                                                  | Level of Output                 |
|--------------------------|------------------------------------------------------------------------------------|---------------------------------|
| <b>Logical Interface</b> |                                                                                    |                                 |
| <b>Logical interface</b> | Name of the logical interface.                                                     | All levels                      |
| <b>Index</b>             | Index number of the logical interface, which reflects its initialization sequence. | <b>detail extensive</b><br>none |
| <b>State</b>             | State of the logical interface: <b>up</b> or <b>down</b> .                         | All levels                      |

Table 174: show pppoe interfaces Output Fields (Continued)

| Field Name                              | Field Description                                                                                                                         | Level of Output                 |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Session ID</b>                       | Session ID.                                                                                                                               | All levels                      |
| <b>Type</b>                             | Origin of the logical interface: <b>Static</b> or <b>Dynamic</b> . Indicates whether the interface was statically or dynamically created. | <b>detail extensive</b><br>none |
| <b>Service name</b>                     | Type of service required (can be used to indicate an ISP name or a class or quality of service).                                          | <b>detail extensive</b><br>none |
| <b>Configured AC name</b>               | Configured access concentrator name.                                                                                                      | <b>detail extensive</b><br>none |
| <b>Session AC name</b>                  | Name of the access concentrator.                                                                                                          | <b>detail extensive</b><br>none |
| <b>Remote MAC address or Remote MAC</b> | MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.                                     | All levels                      |
| <b>Session uptime</b>                   | Length of time the session has been up, in <i>hh:mm:ss</i> .                                                                              | <b>detail extensive</b><br>none |
| <b>Dynamic Profile</b>                  | Name of the dynamic profile that was used to create this interface. If the interface was statically created, this field is not displayed. | <b>detail extensive</b><br>none |
| <b>Underlying interface</b>             | Interface on which PPPoE is running.                                                                                                      | All levels                      |

Table 174: show pppoe interfaces Output Fields (Continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                      | Level of Output                 |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Agent Circuit ID</b>  | Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. If the agent circuit ID is not configured, this field is not displayed.                    | <b>detail extensive</b><br>none |
| <b>Agent Remote ID</b>   | Agent remote identifier that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. If the agent remote ID is not configured, this field is not displayed. | <b>detail extensive</b><br>none |
| <b>ACI Interface Set</b> | Internally-generated name of the dynamic ACI interface set, if configured, and the set index number of the ACI entry in the session database.                                                                                                                                                                                                          | <b>detail extensive</b><br>none |

Table 174: show pppoe interfaces Output Fields (Continued)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Packet Type</b> | <p>Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <ul style="list-style-type: none"> <li>• <b>PADI</b>—PPPoE Active Discovery Initiation packets.</li> <li>• <b>PADO</b>—PPPoE Active Discovery Offer packets.</li> <li>• <b>PADR</b>—PPPoE Active Discovery Request packets.</li> <li>• <b>PADS</b>—PPPoE Active Discovery Session-Confirmation packets.</li> <li>• <b>PADT</b>—PPPoE Active Discovery Termination packets.</li> <li>• <b>Service name error</b>—Packets for which the Service-Name request could not be honored.</li> <li>• <b>AC system error</b>—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit.</li> <li>• <b>Generic error</b>—Packets that indicate an unrecoverable error occurred.</li> <li>• <b>Malformed packets</b>—Malformed or short packets that caused the packet handler to discard the frame as unreadable.</li> <li>• <b>Unknown packets</b>—Unrecognized packets.</li> </ul> | <b>extensive</b> |

## Sample Output

### show pppoe interfaces

```

user@host> show pppoe interfaces
pp0.0 Index 66
 State: Down, Session ID: None,
 Service name: None, Configured AC name: sapphire,
 Session AC name: None, Remote MAC address: 00:00:5e:00:53:00,

```

```
Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
Underlying interface: at-5/0/0.0 Index 71
```

### show pppoe interfaces (Status for the Specified Interface)

```
user@host> show pppoe interfaces pp0.1073741827
pp0.1073741827 Index 70
 State: Session Up, Session ID: 30, Type: Dynamic,
 Session AC name: velorum,
 Remote MAC address: 00:00:5e:00:53:c1,
 Session uptime: 16:45:46 ago,
 Underlying interface: ge-2/0/3.1 Index 73
 Service name: premium
 Dynamic Profile: PppoeProfile
 Agent Circuit ID: velorum-ge-2/0/3
 Agent Remote ID: westford
```

### show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
```

| Interface      | Underlying interface | State      | Session ID | Remote MAC        |
|----------------|----------------------|------------|------------|-------------------|
| pp0.0          | ge-2/0/3.2           | Session Up | 27         | 00:00:5e:00:53:c1 |
| pp0.1          | ge-2/0/3.2           | Session Up | 28         | 00:00:5e:00:53:c1 |
| pp0.1073741824 | ge-2/0/3.1           | Session Up | 29         | 00:00:5e:00:53:c1 |
| pp0.1073741825 | ge-2/0/3.1           | Session Up | 30         | 00:00:5e:00:53:c1 |
| pp0.1073741826 | ge-2/0/3.1           | Session Up | 31         | 00:00:5e:00:53:c1 |

### show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 66
 State: Down, Session ID: None, Type: Static,
 Service name: None, Configured AC name: sapphire,
 Session AC name: None, Remote MAC address: 00:00:5e:00:53:00,
 Auto-reconnect timeout: 100 seconds, Idle timeout: Never,
 Underlying interface: at-5/0/0.0 Index 71
```

## show pppoe interfaces (PPPoE Subscriber Interface with ACI Interface Set)

```
user@host> show pppoe interfaces pp0.1073741827
pp0.1073741827 Index 346
 State: Session Up, Session ID: 4, Type: Dynamic,
 Service name: AGILENT, Remote MAC address: 00:00:5e:00:53:62,
 Session AC name: nbc,
 Session uptime: 6d 02:22 ago,
 Dynamic Profile: aci-vlan-pppoe-profile,
 Underlying interface: demux0.1073741826 Index 345
 Agent Circuit ID: aci-ppp-dhcp-dvlan-50
 ACI Interface Set: aci-1002-demux0.1073741826 Index 2
```

### Release Information

Command introduced before Junos OS Release 7.4.

### RELATED DOCUMENTATION

| [Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration](#)

## show pppoe service-name-tables

### IN THIS SECTION

- [Syntax | 1772](#)
- [Description | 1772](#)
- [Options | 1772](#)
- [Required Privilege Level | 1772](#)
- [Output Fields | 1772](#)
- [Sample Output | 1774](#)
- [Release Information | 1775](#)

## Syntax

```
show pppoe service-name-tables
<table-name>
```

## Description

Display configuration information about PPPoE service name tables.

## Options

**none** Display the names of configured PPPoE service name tables.

***table-name*** (Optional) Name of a configured PPPoE service name table.

## Required Privilege Level

view

## Output Fields

[Table 175 on page 1772](#) lists the output fields for the **show pppoe service-name-tables** command. Output fields are listed in the approximate order in which they appear.

**Table 175: show pppoe service-name-tables Output Fields**

| Field Name                | Field Description                     | Level of Output |
|---------------------------|---------------------------------------|-----------------|
| <b>Service Name Table</b> | Name of the PPPoE service name table. | none            |

Table 175: show pppoe service-name-tables Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Service Name</b>     | <p>Name of a configured service in the PPPoE service name table:</p> <ul style="list-style-type: none"> <li>• <b>&lt;empty&gt;</b>—Service of zero length that represents an unspecified service</li> <li>• <b>&lt;any&gt;</b>—Default service for non-empty service entries that do not match the configured empty or named service entries</li> <li>• <b>service-name</b>—Named service entry</li> </ul>                                                                                                                                                    | none            |
| <b>Action</b>           | <p>Action taken when the PPPoE underlying interface interface receives a PPPoE Active Discovery Initiation (PADI) packet with the specified named service, <b>empty</b> service, <b>any</b> service, or ACI/ARI pair:</p> <ul style="list-style-type: none"> <li>• <b>Delay seconds</b>—Number of seconds that the interface delays before responding with a PPPoE Active Discovery Offer (PADO) packet</li> <li>• <b>Drop</b>—Interface drops (ignores) the packet.</li> <li>• <b>Terminate</b>—Interface responds immediately with a PADO packet</li> </ul> | none            |
| <b>Dynamic Profile</b>  | <p>Name of the dynamic profile with which the router creates a dynamic PPPoE subscriber interface. A dynamic profile can be assigned to a named service, <b>empty</b> service, <b>any</b> service, or ACI/ARI pair.</p>                                                                                                                                                                                                                                                                                                                                       | none            |
| <b>Routing Instance</b> | <p>Name of the routing instance in which to instantiate the dynamic PPPoE subscriber interface. A routing instance can be assigned to a named service, <b>empty</b> service, <b>any</b> service, or ACI/ARI pair.</p>                                                                                                                                                                                                                                                                                                                                         | none            |
| <b>Max Sessions</b>     | <p>Maximum number of active PPPoE sessions that the router can establish with the specified named service, <b>empty</b> service, or <b>any</b> service.</p>                                                                                                                                                                                                                                                                                                                                                                                                   | none            |

Table 175: show pppoe service-name-tables Output Fields (Continued)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Active Sessions</b>  | Current count of active PPPoE sessions created using the specified named service, <b>empty</b> service, or <b>any</b> service. The Active Sessions value cannot exceed the Max Sessions value.                                                                                                                                                                                                              | none            |
| <b>ACI</b>              | Agent circuit identifier (ACI) that corresponds to the DSLAM interface that initiated the client service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both the beginning and end of the string. An ACI can be configured as part of an ACI/ARI pair for a named service, <b>empty</b> service, or <b>any</b> service.                           | none            |
| <b>ARI</b>              | Agent remote identifier (ARI) that corresponds to the subscriber associated with the DSLAM interface that initiated the service request. An asterisk is interpreted as a wildcard character and can appear at the beginning, the end, or both at the beginning and end of the string. An ARI can be configured as part of an ACI/ARI pair for a named service, <b>empty</b> service, or <b>any</b> service. | none            |
| <b>Static Interface</b> | Name of the static PPPoE interface reserved for exclusive use by the PPPoE client with matching ACI/ARI information. A static interface can be configured only for an ACI/ARI pair.                                                                                                                                                                                                                         | none            |

## Sample Output

### show pppoe service-name-tables

```

user@host> show pppoe service-name-tables
Service Name Table: test1
Service Name Table: test2
Service Name Table: test3

```

## show pppoe service-name-tables (For the Specified Table Name)

```
user@host> show pppoe service-name-tables Table1
Service Name Table: Table1
 Service Name: <empty>
 Action: Terminate
 Dynamic Profile: BasicPppoeProfile
 Max Sessions: 100
 Active Sessions: 3
 Service Name: <any>
 Action: Drop
 ACI: velorum-ge-2/0/3
 ARI: westford
 Action: Terminate
 Static Interface: pp0.100
 ACI: volantis-ge-5/0/5
 ARI: sunnyvale
 Action: Terminate
 Static Interface: pp0.101
 Service Name: Wholesale
 Action: Terminate
 Dynamic Profile: WholesalePppoeProfile
 Routing Instance: WholesaleRI
 Max Sessions: 16000
 Active Sessions: 4
```

## Release Information

Command introduced in Junos OS Release 10.0.

## RELATED DOCUMENTATION

[Verifying a PPPoE Configuration](#)

*Verifying and Managing Dynamic PPPoE Configuration*

## show pppoe sessions

### IN THIS SECTION

- [Syntax | 1776](#)
- [Description | 1776](#)
- [Options | 1776](#)
- [Required Privilege Level | 1777](#)
- [Output Fields | 1777](#)
- [Sample Output | 1777](#)
- [Release Information | 1778](#)

### Syntax

```
show pppoe sessions
<aci circuit-id-string>
<ari remote-id-string>
<service service-name>
```

### Description

Display information about all active PPPoE sessions on the router, or about the active PPPoE sessions established for a specified service name, agent circuit identifier (ACI), or agent remote identifier (ARI).

### Options

- |                                     |                                                                                                                                                                                                                                                         |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>none</b>                         | Display information for all active PPPoE sessions on the router.                                                                                                                                                                                        |
| <b>aci <i>circuit-id-string</i></b> | (Optional) Display information only for active PPPoE sessions established with the specified agent circuit identifier. The agent circuit identifier corresponds to the DSLAM interface that initiated the service request.                              |
| <b>ari <i>remote-id-string</i></b>  | (Optional) Display information only for active PPPoE sessions established with the specified agent remote identifier. The agent remote identifier corresponds to the subscriber associated with the DSLAM interface that initiated the service request. |

**service *service-name*** (Optional) Display information only for active PPPoE sessions established with the specified service, where *service-name* can be **empty**, **any**, or a named service.

## Required Privilege Level

view

## Output Fields

[Table 176 on page 1777](#) lists the output fields for the **show pppoe sessions** command. Output fields are listed in the approximate order in which they appear.

**Table 176: show pppoe sessions Output Fields**

| Field Name                  | Field Description                                                                                     | Level of Output |
|-----------------------------|-------------------------------------------------------------------------------------------------------|-----------------|
| <b>Interface</b>            | Name of the statically-created or dynamically-created PPPoE interface for the active PPPoE session.   | none            |
| <b>Underlying interface</b> | Interface on which PPPoE is running.                                                                  | none            |
| <b>State</b>                | State of the PPPoE session; displays <b>Session Up</b> for active PPPoE sessions.                     | none            |
| <b>Session ID</b>           | PPPoE session identifier.                                                                             | none            |
| <b>Remote MAC</b>           | MAC address of the remote side of the connection, either the access concentrator or the PPPoE client. | none            |

## Sample Output

### show pppoe sessions (For All Active Sessions)

```
user@host> show pppoe sessions
Interface Underlying State Session Remote
```

|                | interface  |            | ID | MAC               |
|----------------|------------|------------|----|-------------------|
| pp0.0          | ge-2/0/3.2 | Session Up | 27 | 00:00:5e:00:53:c1 |
| pp0.1          | ge-2/0/3.2 | Session Up | 28 | 00:00:5e:00:53:c1 |
| pp0.1073741824 | ge-2/0/3.1 | Session Up | 29 | 00:00:5e:00:53:c1 |
| pp0.1073741825 | ge-2/0/3.1 | Session Up | 30 | 00:00:5e:00:53:c1 |
| pp0.1073741826 | ge-2/0/3.1 | Session Up | 31 | 00:00:5e:00:53:c1 |

### show pppoe sessions (For All Active Sessions Matching the Agent Circuit Identifier)

```

user@host> show pppoe sessions aci "velorum-ge-2/0/3"
Interface Underlying State Session Remote
 interface
 ID MAC
pp0.0 ge-2/0/3.2 Session Up 27 00:00:5e:00:53:c1
pp0.1 ge-2/0/3.2 Session Up 28 00:00:5e:00:53:c1

```

## Release Information

Command introduced in Junos OS Release 10.2.

## RELATED DOCUMENTATION

[Verifying a PPPoE Configuration](#)

*Verifying and Managing Dynamic PPPoE Configuration*

## show pppoe statistics

### IN THIS SECTION

- [Syntax | 1779](#)
- [Description | 1779](#)
- [Options | 1779](#)
- [Required Privilege Level | 1779](#)
- [Output Fields | 1779](#)

- [Sample Output | 1781](#)
- [Release Information | 1782](#)

## Syntax

```
show pppoe statistics
<logical-interface-name>
```

## Description

Display statistics information about PPPoE interfaces.

## Options

- |                                      |                                                          |
|--------------------------------------|----------------------------------------------------------|
| <b>none</b>                          | Display PPPoE statistics for all interfaces.             |
| <b><i>logical-interface-name</i></b> | (Optional) Name of a PPPoE underlying logical interface. |

## Required Privilege Level

view

## Output Fields

[Table 177 on page 1780](#) lists the output fields for the **show pppoe statistics** command. Output fields are listed in the approximate order in which they appear.

Table 177: show pppoe statistics Output Fields

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Active PPPoE sessions</b> | <p>Total number of active PPPoE sessions and the number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <ul style="list-style-type: none"> <li>• <b>PADI</b>—PPPoE Active Discovery Initiation packets.</li> <li>• <b>PADO</b>—PPPoE Active Discovery Offer packets.</li> <li>• <b>PADR</b>—PPPoE Active Discovery Request packets.</li> <li>• <b>PADS</b>—PPPoE Active Discovery Session-Confirmation packets.</li> <li>• <b>PADT</b>—PPPoE Active Discovery Termination packets.</li> <li>• <b>Service name error</b>—Packets for which the Service-Name request could not be honored.</li> <li>• <b>AC system error</b>—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit.</li> <li>• <b>Generic error</b>—Packets that indicate an unrecoverable error occurred.</li> <li>• <b>Malformed packets</b>—Malformed or short packets that caused the packet handler to discard the frame as unreadable.</li> <li>• <b>Unknown packets</b>—Unrecognized packets.</li> </ul> |
| <b>Timeouts</b>              | <p>Information about timeouts that occurred during the PPPoE session (not displayed for M120, M320, and MX Series routers):</p> <ul style="list-style-type: none"> <li>• <b>PADI</b>—No PADR packet has been received within the timeout period. (This value is always zero and is not supported.)</li> <li>• <b>PADO</b>—No PPPoE Active Discovery Offer packet has been received within the timeout period.</li> <li>• <b>PADR</b>—No PADS packet has been received within the timeout period.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

### show pppoe statistics

```
user@host> show pppoe statistics
Active PPPoE sessions: 1
 PacketType Sent Received
 PADI 0 0
 PADO 0 0
 PADR 0 0
 PADS 0 0
 PADT 0 0
 Service name error 0 0
 AC system error 0 0
 Generic error 0 0
 Malformed packets 0 0
 Unknown packets 0 0
 Timeouts
 PADI 0
 PADO 0
 PADR 0
```

### show pppoe statistics (For the Specified Underlying Interface Only)

```
user@host> show pppoe statistics ge-4/0/3.2
Active PPPoE sessions: 4
 PacketType Sent Received
 PADI 0 5
 PADO 5 0
 PADR 0 5
 PADS 4 0
 PADT 0 1
 Service name error 0 0
 AC system error 0 0
 Generic error 0 0
 Malformed packets 0 0
 Unknown packets 0 0
```

## Release Information

Command introduced before Junos OS Release 7.4.

*logical-interface-name* option introduced in Junos OS Release 10.1.

## RELATED DOCUMENTATION

[show ppp address-pool](#)

[show pppoe underlying-interfaces](#) | 1782

## show pppoe underlying-interfaces

### IN THIS SECTION

- [Syntax](#) | 1782
- [Description](#) | 1782
- [Options](#) | 1783
- [Required Privilege Level](#) | 1783
- [Output Fields](#) | 1783
- [Sample Output](#) | 1789
- [Release Information](#) | 1792

## Syntax

```
show pppoe underlying-interfaces
<brief | detail | extensive>
<lockout>
<logical-interface-name>
```

## Description

Display information about PPPoE underlying interfaces.

## Options

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

**lockout** (Optional) Display summary information about the lockout condition and the lockout grace period for PPPoE clients on the PPPoE underlying interface.

***logical-interface-name*** (Optional) Name of a PPPoE underlying logical interface.

## Required Privilege Level

view

## Output Fields

[Table 178 on page 1783](#) lists the output fields for the **show pppoe underlying-interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 178: show pppoe underlying-interfaces Output Fields**

| Field Name                  | Field Description                                                                                                                            | Level of Output         |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Underlying Interface</b> | Name of the PPPoE underlying logical interface.                                                                                              | All levels              |
| <b>Service Name Table</b>   | Name of the service name table.                                                                                                              | All levels              |
| <b>Dynamic Profile</b>      | Name of the dynamic profile that was used to create this interface. If the interface was statically created, then the value is <b>none</b> . | All levels              |
| <b>Index</b>                | Index number of the logical interface, which reflects its initialization sequence.                                                           | <b>detail extensive</b> |
| <b>State</b>                | Origin of the logical interface: <b>Static</b> or <b>Dynamic</b> . Indicates whether the interface was statically or dynamically created.    | <b>detail extensive</b> |

Table 178: show pppoe underlying-interfaces Output Fields (Continued)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                      | Level of Output                 |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Operational States</b>      | Fields in this block are actual operational values rather than simply the configured values. The operational values can be the result of RADIUS-initiated changes.                                                                                                                                                                                                     | <b>detail extensive</b>         |
| <b>Max Sessions</b>            | Maximum number of PPPoE logical interfaces that can be activated on the underlying interface. When this number of logical interfaces has been established, all subsequent PPPoE Active Discovery Initiation (PADI) packets are dropped and all subsequent PPPoE Active Discovery Request (PADR) packets trigger PPPoE Active Discovery Session (PADS) error responses. | <b>detail extensive</b>         |
| <b>Max Sessions VSA Ignore</b> | Whether the router is configured to ignore (clear) the PPPoE maximum session value returned by RADIUS in the Max-Clients-Per-Interface Juniper Networks VSA [26-143] and restore the PPPoE maximum session value on the underlying interface to the value configure with the <b>max-sessions</b> statement: <b>Off</b> (default) or <b>On</b> .                        | <b>detail extensive</b><br>none |
| <b>Active Sessions</b>         | Number of active PPPoE sessions on the underlying interface. If a dynamic profile is listed, then it is the number of active PPPoE sessions on the underlying interface that are using this profile. The Active Sessions value must not exceed the Max Sessions value.                                                                                                 | <b>detail extensive</b>         |

Table 178: show pppoe underlying-interfaces Output Fields (Continued)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Agent Circuit Identifier</b> | <p>Whether the underlying interface is configured with the <b>agent-circuit-identifier</b> statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on agent circuit identifier (ACI) information.</p> <p><b>Autosensing</b> indicates that creation of ACI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ACI-based dynamic VLANs is not configured on the underlying interface, this field does not appear.</p> <p><b>NOTE:</b> The Agent Circuit Identifier field is replaced with the Line Identity field when an ALI interface set is configured with the <b>line-identity</b> autoconfiguration stanza.</p>          | <b>detail extensive</b><br>none |
| <b>Line Identity</b>            | <p>Whether the underlying interface is configured with the <b>line-identity</b> statement to enable creation of autosensed dynamic VLAN subscriber interfaces based on the specified trusted option: ACI, ARI, both, or neither.</p> <p><b>Autosensing</b> indicates that creation of ALI-based dynamic VLAN interfaces is enabled on the underlying interface. If creation of ALI dynamic VLANs based on trusted options is not configured on the underlying interface, this field does not appear.</p> <p><b>NOTE:</b> The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the <b>agent-circuit-id</b> autoconfiguration stanza.</p> | <b>detail extensive</b><br>none |
| <b>Duplicate Protection</b>     | <p>State of PPPoE duplicate protection: <b>On</b> or <b>Off</b>. When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client. The uniqueness of the PPPoE client is determined by the client's MAC address.</p>                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b>         |

Table 178: show pppoe underlying-interfaces Output Fields (Continued)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Level of Output                 |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Short Cycle Protection</b> | State of PPPoE short cycle protection: <b>mac-address</b> , <b>circuit-id</b> , or <b>Off</b> . Enabling short cycle protection, also known as PPPoE lockout, on the PPPoE underlying interface temporarily prevents (locks out) a failed or short-lived (short-cycle) PPPoE subscriber session from reconnecting to the router for a default or configurable period of time. PPPoE client sessions are identified by their unique media access control (MAC) source address or agent circuit identifier (ACI) value. | <b>detail extensive</b>         |
| <b>Direct Connect</b>         | State of the configuration to ignore DSL Forum VSAs: <b>On</b> or <b>Off</b> . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.                                                                                                                                                                                                                                                                                                                  | <b>detail extensive</b><br>none |
| <b>AC Name</b>                | Name of the access concentrator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>detail extensive</b>         |

Table 178: show pppoe underlying-interfaces Output Fields (*Continued*)

| Field Name        | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>PacketType</b> | <p>Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <ul style="list-style-type: none"> <li>• <b>PADI</b>—PPPoE Active Discovery Initiation packets.</li> <li>• <b>PADO</b>—PPPoE Active Discovery Offer packets.</li> <li>• <b>PADR</b>—PPPoE Active Discovery Request packets.</li> <li>• <b>PADS</b>—PPPoE Active Discovery Session-Confirmation packets.</li> <li>• <b>PADT</b>—PPPoE Active Discovery Termination packets.</li> <li>• <b>Service name error</b>—Packets for which the Service-Name request could not be honored.</li> <li>• <b>AC system error</b>—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit.</li> <li>• <b>Generic error</b>—Packets that indicate an unrecoverable error occurred.</li> <li>• <b>Malformed packets</b>—Malformed or short packets that caused the packet handler to discard the frame as unreadable.</li> <li>• <b>Unknown packets</b>—Unrecognized packets.</li> </ul> | <b>detail extensive</b> |

Table 178: show pppoe underlying-interfaces Output Fields (Continued)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output  |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Lockout Time (sec)</b> | <p>The PPPoE lockout time range, the number of PPPoE clients in lockout condition, and the number of PPPoE clients in a lockout grace period if <b>Short Cycle Protection</b> is enabled (<b>On</b>):</p> <ul style="list-style-type: none"> <li>• <b>Min</b>—Minimum lockout time, in seconds, configured on the PPPoE underlying interface.</li> <li>• <b>Max</b>—Maximum lockout time, in seconds, configured on the PPPoE underlying interface.</li> <li>• <b>Total clients in lockout</b>—Number of PPPoE clients currently undergoing lockout.</li> <li>• <b>Total clients in lockout grace period</b>—Number of PPPoE clients currently in a lockout grace period. A <i>lockout grace period</i> occurs when the time between lockout events is greater than either 15 minutes or the maximum lockout time.</li> </ul> | <b>extensive</b> |
| <b>Client Address</b>     | MAC source address of the PPPoE client.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b> |
| <b>Current</b>            | Current lockout time, in seconds; displays <b>0</b> (zero) if the PPPoE client is not undergoing lockout.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>extensive</b> |
| <b>Elapsed</b>            | Time elapsed into the lockout period, in seconds; displays <b>0</b> if the PPPoE client is not undergoing lockout                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>extensive</b> |
| <b>Next</b>               | Lockout time, in seconds, that the router uses for the next lockout event; displays a nonzero value if the PPPoE client is currently in a lockout grace period.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>extensive</b> |

## Sample Output

### show pppoe underlying-interfaces brief

```
user@host> show pppoe underlying-interfaces brief
Underlying Interface Service Name Table Dynamic Profile
ge-4/0/3.1 Premium None
ge-4/0/3.2 None PppoeProfile
```

### show pppoe underlying-interfaces detail

```
user@host> show pppoe underlying-interfaces detail
ge-4/0/3.1 Index 73
 Operational States:
 State: Static, Dynamic Profile: None,
 Max Sessions: 4000, Max Sessions VSA Ignore: Off,
 Active Sessions: 0,
 Service Name Table: Premium,
 Direct Connect: Off,
 AC Name: velorum, Duplicate Protection: On,
 Short Cycle Protection: Off

ge-4/0/3.2 Index 78
 Operational States:
 State: Dynamic, Dynamic Profile: PppoeProfile,
 Max Sessions: 500, Max Sessions VSA Ignore: Off,
 Active Sessions: 3,
 Service Name Table: None,
 Direct Connect: Off,
 AC Name: velorum, Duplicate Protection: On,
 Short Cycle Protection: Off
```

### show pppoe underlying-interfaces extensive

```
user@host> show pppoe underlying-interfaces extensive
ge-4/0/3.1 Index 73
 Operational States:
```

State: Static, Dynamic Profile: None,  
 Max Sessions: 4000, Max Sessions VSA Ignore Off,  
 Active Sessions: 0,  
 Service Name Table: None,  
 Direct Connect: Off,  
 AC Name: velorum, Duplicate Protection: Off,  
 Short Cycle Protection: Off

| PacketType         | Sent | Received |
|--------------------|------|----------|
| PADI               | 0    | 0        |
| PADO               | 0    | 0        |
| PADR               | 0    | 0        |
| PADS               | 0    | 0        |
| PADT               | 0    | 0        |
| Service name error | 0    | 0        |
| AC system error    | 0    | 0        |
| Generic error      | 0    | 0        |
| Malformed packets  | 0    | 0        |
| Unknown packets    | 0    | 0        |

ge-4/0/3.2 Index 78

Operational States:

State: Dynamic, Dynamic Profile: PppoeProfile,  
 Max Sessions: 4000, Max Sessions VSA Ignore: Off  
 Active Sessions: 3,  
 Service Name Table: None,  
 Direct Connect: Off,  
 AC Name: velorum, Duplicate Protection: Off,  
 Short Cycle Protection: Off

| PacketType         | Sent | Received |
|--------------------|------|----------|
| PADI               | 0    | 5        |
| PADO               | 5    | 0        |
| PADR               | 0    | 5        |
| PADS               | 4    | 0        |
| PADT               | 0    | 1        |
| Service name error | 0    | 0        |
| AC system error    | 0    | 0        |
| Generic error      | 0    | 0        |
| Malformed packets  | 0    | 0        |
| Unknown packets    | 0    | 0        |

**show pppoe underlying-interfaces extensive (PPPoE client in lockout condition)**

```

user@host> show pppoe underlying-interfaces ge-1/0/0.0 extensive
ge-1/0/0.0 Index 71
 State: Static, Dynamic Profile: None,
 Max Sessions: 32000, Max Sessions VSA Ignore: Off,
 Active Sessions: 0,
 Service Name Table: None,
 Direct Connect: Off,
 AC name: winona, Duplicate Protection: On,
 Short Cycle Protection: Off

PacketType Sent Received

PADI 0 7
PADO 3 0
PADR 0 3
PADS 3 0
PADT 2 1
Service name error 0 0
AC system error 0 0
Generic error 0 0
Malformed packets 0 0
Unknown packets 0 0

Lockout Time (sec): Min: 1, Max: 30
Total clients in lockout: 1
Total clients in lockout grace period: 0

Client Address Current Elapsed Next

00:00:5e:00:53:11 4 3 8

```

**show pppoe underlying-interfaces lockout**

```

user@host> show pppoe underlying-interfaces ge-1/0/0.0 lockout
ge-1/0/0.0 Index 71
 Short Cycle Protection: Off,
 Lockout Time (sec): Min: 10, Max: 60
 Total clients in lockout: 0
 Total clients in lockout grace period: 0

```

## show pppoe underlying-interfaces detail (Autosensing Configured for ACI-based Dynamic VLANs)

```
user@host> show pppoe underlying-interfaces demux0.1073741826 detail
demux0.1073741826 Index 345
 State: Dynamic, Dynamic Profile: aci-vlan-pppoe-profile,
 Max Sessions: 32000, Max Sessions VSA Ignore: Off,
 Active Sessions: 1,
 Agent Circuit Identifier: Autosensing,
 Service Name Table: None,
 Duplicate Protection: On, Short Cycle Protection: Off,
 Direct Connect: Off,
 AC Name: nbc,
 Short Cycle Protection: circuit-id,
```

## show pppoe underlying-interfaces detail (Autosensing Configured for ALI-based Dynamic VLANs)

```
user@host> show pppoe underlying-interfaces demux0.1073741826 detail
demux0.1073741826 Index 345
 State: Dynamic, Dynamic Profile: aci-vlan-pppoe-profile,
 Max Sessions: 32000, Max Sessions VSA Ignore: Off,
 Active Sessions: 1,
 Line Identity: Autosensing,
 Service Name Table: None,
 Duplicate Protection: On, Short Cycle Protection: Off,
 Direct Connect: Off,
 AC Name: nbc,
 Short Cycle Protection: circuit-id,
```

## Release Information

Command introduced in Junos OS Release 10.0.

**lockout** option added in Junos OS Release 11.4.

## RELATED DOCUMENTATION

*Verifying and Managing Dynamic PPPoE Configuration*

*Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces*

*Configuring the PPPoE Family for an Underlying Interface*

*Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

*Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

## show pppoe version

### IN THIS SECTION

- [Syntax | 1793](#)
- [Description | 1793](#)
- [Options | 1793](#)
- [Required Privilege Level | 1794](#)
- [Output Fields | 1794](#)
- [Sample Output | 1795](#)
- [Release Information | 1795](#)

## Syntax

```
show pppoe version
```

## Description

( M120 routers and M320 routers only) Display version information about PPPoE.

## Options

This command has no options.

## Required Privilege Level

view

## Output Fields

Table 179 on page 1794 lists the output fields for the **show pppoe version** command. Output fields are listed in the approximate order in which they appear.

**Table 179: show pppoe version Output Fields**

| Field Name                       | Field Description                                                                                                                                                                                                                                                                   |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>version <i>n</i></b>          | PPPoE version number and RFC. For example, <b>version 1, rfc 2516</b> .                                                                                                                                                                                                             |
| <b>PPPoE protocol</b>            | State of the PPPoE protocol: <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                    |
| <b>Maximum Sessions</b>          | Maximum active sessions supported per router. The default is 256 sessions.                                                                                                                                                                                                          |
| <b>PADI resend timeout</b>       | Initial time, in seconds, that the router waits to receive a PPOE Active Discovery Offer (PADO) packet for the PPOE Active Discovery Initiation (PADI) packet sent. This timeout doubles for each successive PADI packet sent. Not displayed for M120 and M320 routers.             |
| <b>PADR resend timeout</b>       | Initial time, in seconds, that the router waits to receive a PPOE Active Discovery Session Confirmation (PADS) packet for the PPOE Active Discovery Request (PADR) packet sent. This timeout doubles for each successive PADR packet sent. Not displayed for M120 and M320 routers. |
| <b>Max resend timeout</b>        | Maximum value, in seconds, that the PADI or PADR resend timer can accept. The maximum value is 64. Not displayed for M120 and M320 routers.                                                                                                                                         |
| <b>Max Configured AC timeout</b> | Time, in seconds, during which the configured access concentrator must respond. Not displayed for M120 and M320 routers.                                                                                                                                                            |

## Sample Output

### show pppoe version

```
user@host> show pppoe version
Point-to-Point Protocol Over Ethernet, version 1. rfc2516
 PPPoE protocol = Enabled
 Maximum Sessions = 256
 PADI resend timeout = 2 seconds
 PADR resend timeout = 16 seconds
 Max resend timeout = 64 seconds
 Max Configured AC timeout = 4 seconds
```

## Release Information

Command introduced before Junos OS Release 7.4.

## traceroute ethernet

### IN THIS SECTION

- [Syntax | 1796](#)
- [Description | 1796](#)
- [Options | 1796](#)
- [Required Privilege Level | 1797](#)
- [Output Fields | 1797](#)
- [Sample Output | 1799](#)
- [Release Information | 1800](#)

## Syntax

```

traceroute ethernet
local-mep mep-id
maintenance-association ma-name
maintenance-domain md-name
<tll value>
<wait seconds>
mac-address | mep-id
<detail>

```

## Description

Triggers the linktrace protocol to trace the route between two maintenance points. The result of the traceroute protocol is stored in the path database. To display the path database, use the **show oam ethernet connectivity-fault-management path-database** command.

Before using the traceroute command, you can verify the remote MEP's MAC address using the **show oam ethernet connectivity-fault-management path-database** command.

## Options

|                                               |                                                                                                                                |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>local-mep</b> <i>mep-id</i>                | (Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.                                    |
| <b>detail</b>                                 | (Optional) Provide detailed information of the responder hostname, ingress port name, egress port name, TTL, and relay action. |
| <b>mac-address</b>                            | Destination unicast MAC address of the remote maintenance point.                                                               |
| <b><i>mep-id</i></b>                          | MEP identifier of the remote maintenance point. The range of values is 1 through 8191.                                         |
| <b>maintenance-association</b> <i>ma-name</i> | Specifies an existing maintenance association from the set of configured maintenance associations.                             |
| <b>maintenance-domain</b> <i>md-name</i>      | Specifies an existing maintenance domain from the set of configured maintenance domains.                                       |
| <b>tll</b> <i>value</i>                       | Number of hops to use in the linktrace request. The range is 1 to 255 hops. The default is 4.                                  |

**wait *seconds*** (Optional) Maximum time to wait for a response to the traceroute request. The range is 1 to 255 seconds. The default is 5.

## Required Privilege Level

network

## Output Fields

[Table 180 on page 1797](#) lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

**Table 180: traceroute ethernet Output Fields**

| Field Name                     | Field Description                                                                                                                                                                                                                                                            |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Linktrace to</b>            | MAC address of the destination maintenance point.                                                                                                                                                                                                                            |
| <b>Interface</b>               | Local interface used to send the linktrace message (LTM).                                                                                                                                                                                                                    |
| <b>Maintenance Domain</b>      | Maintenance domain specified in the traceroute command.                                                                                                                                                                                                                      |
| <b>Level</b>                   | Maintenance domain level configured.                                                                                                                                                                                                                                         |
| <b>Maintenance Association</b> | Maintenance association specified in the traceroute command.                                                                                                                                                                                                                 |
| <b>Local Mep</b>               | The local maintenance end point identifier.                                                                                                                                                                                                                                  |
| <b>Transaction Identifier</b>  | 4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all Maintenance Domains. Use the transaction identifier to match an incoming linktrace response (LTR), with a previously sent LTM. |
| <b>Hop</b>                     | Sequential hop count of the linktrace path.                                                                                                                                                                                                                                  |

Table 180: traceroute ethernet Output Fields *(Continued)*

| Field Name                  | Field Description                                                                                                                                                                                                                                                  |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>TTL</b>                  | Number of hops remaining in the linktrace message. The time to live (TTL) is decremented at each hop.                                                                                                                                                              |
| <b>Source MAC address</b>   | MAC address of the 802.1ag node responding to the LTM or the source MAC address of the LTR.                                                                                                                                                                        |
| <b>Next-hop MAC address</b> | MAC address of the egress interface of the node to which the LTM is forwarded or<br><br>the next-hop MAC address derived from the next egress identifier in the Egress-ID TLV of the LTR PDU.                                                                      |
| <b>Responder Hostname</b>   | The hostname of the responding router. A valid hostname is received only when the responding system is a Juniper Networks router.                                                                                                                                  |
| <b>Ingress port name</b>    | The port name for ingress connections.                                                                                                                                                                                                                             |
| <b>Egress port name</b>     | The port name for egress connections.                                                                                                                                                                                                                              |
| <b>Flags</b>                | The configurable flags can include: <ul style="list-style-type: none"> <li>• H– Hardware only, incoming LT frame has hardware bit set.</li> <li>• T– Terminal MEP, responder is a terminating MEP.</li> <li>• F– FWD yes, LTM frame is relayed further.</li> </ul> |

Table 180: traceroute ethernet Output Fields (Continued)

| Field Name          | Field Description                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Relay Action</b> | <p>The associated relay action. Relay action can be one of the following:</p> <ul style="list-style-type: none"> <li>• RlyHit— Relay hit; target MAC address matches the MP mac address.</li> <li>• RlyFDB— Relay FDB; output port decided by consulting forwarding database.</li> <li>• RlyMPDB— Relay MIP; output port decided by consulting MIP database.</li> </ul> |

## Sample Output

### traceroute ethernet

```
user@host> traceroute ethernet maintenance-domain md1 maintenance-association ma1
```

```
00:01:02:03:04:05
```

```
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
```

```
Maintenance Domain: MD1, Level: 7
```

```
Maintenance Association: MA1, Local Mep: 1
```

| Hop                           | TTL | Source MAC address | Next hop MAC address |
|-------------------------------|-----|--------------------|----------------------|
| Transaction Identifier:100001 |     |                    |                      |
| 1                             | 63  | 00:00:aa:aa:aa:aa  | 00:00:ab:ab:ab:ab    |
| 2                             | 62  | 00:00:bb:bb:bb:bb  | 00:00:bc:bc:bc:bc    |
| 3                             | 61  | 00:00:cc:cc:cc:cc  | 00:00:cd:cd:cd:cd    |
| 4                             | 60  | 00:01:02:03:04:05  | 00:00:00:00:00:00    |

### traceroute ethernet detail

```
user@host> run traceroute ethernet maintenance-domain md6 maintenance-association ma6 mep 101
```

```
detail
```

```
Linktrace to 00:00:5E:00:53:CC, Interface : ge-1/0/0.1
```

```
Maintenance Domain: md6, Level: 6
```

```
Maintenance Association: ma6, Local Mep: 201
```

Transaction Identifier: 2077547465

Legend for RelayAction:

RlyHit -- Relay hit, Target MAC address matches the MP mac address  
 RlyFDB -- Relay FDB, output port decided by consulting FDB database  
 RlyMPDB -- Relay MIP, output port decided by consulting MIP database

Legend for Flags:

H -- Hardware only, incoming LT frame has hardware bit set  
 T -- Terminal MEP, responder is a terminating MEP  
 F -- FWD yes, LTM frame is relayed further

| TTL | Responder Hostname | Ingress port name   | Egress port name   | RelayAction |
|-----|--------------------|---------------------|--------------------|-------------|
|     | Responder Service  | Ingress MAC address | Egress MAC address | Flags       |
| 62  | host1              | ge-1/0/0.1          | ge-2/3/0.1         | RlyFDB      |
|     | br1                | 00:00:5E:00:53:00   | 00:00:5E:00:53:A0  | HF-         |
| 63  | host2              | ge-2/3/0.1          | ge-1/0/0.1         | RlyFDB      |
|     | br1                | 00:00:5E:00:53:AA   | 00:00:5E:00:53:A2  | HF-         |
| 61  | host3              | ge-1/0/0.1          | --:--              | RlyHit      |
|     | br1                | 00:00:5E:00:53:B0   | --:--              | H-T         |

## Release Information

Command introduced in Junos OS Release 9.0.

**mep-id** option introduced in Junos OS Release 9.1.

**local-mep** option introduced in Junos OS Release 15.1