

**Junos<sup>®</sup> OS**

---

# Ethernet Interfaces User Guide for Routing Devices

Published  
2019-12-26

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://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<sup>®</sup> OS Ethernet Interfaces User Guide for Routing Devices*  
Copyright © 2019 Juniper Networks, Inc. All rights reserved.

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

## YEAR 2000 NOTICE

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

## END USER LICENSE AGREEMENT

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

# Table of Contents

## About the Documentation | xxxvii

Documentation and Release Notes | xxxvii

Using the Examples in This Manual | xxxvii

    Merging a Full Example | xxxviii

    Merging a Snippet | xxxix

Documentation Conventions | xxxix

Documentation Feedback | xlii

Requesting Technical Support | xlii

    Self-Help Online Tools and Resources | xliii

    Creating a Service Request with JTAC | xliii

## 1

## Ethernet Interfaces

### Configuring Ethernet Interfaces | 3

Ethernet Interfaces Overview | 3

Initial Configuration of Ethernet Interfaces | 4

    Configuring Ethernet Physical Interface Properties | 5

    Configuring the Interface Speed on Ethernet Interfaces | 7

    Configuring the Ingress Rate Limit | 8

    Configuring the Link Characteristics on Ethernet Interfaces | 9

    Configuring Multicast Statistics Collection on Ethernet Interfaces | 10

    MAC Address Validation on Static Ethernet Interfaces Overview | 11

    Configuring MAC Address Validation on Static Ethernet Interfaces | 12

    Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router | 13

    Example: Configuring Fast Ethernet Interfaces | 15

    Example: Configuring Gigabit Ethernet Interfaces | 16

Flow Control for Ethernet Interfaces | 17

    Understanding Flow Control | 18

        IEEE 802.3X Ethernet PAUSE | 18

    Configuring Flow Control | 19

MAC Address Filtering and Accounting on Ethernet Interfaces	20
Configuring MAC Address Filtering for Ethernet Interfaces	21
Enabling Source Address Filtering	21
Configuring MAC Address Filtering on PTX Series Packet Transport Routers	23
Configuring MAC Address Accounting	24
Management Ethernet Interfaces	25
Management Ethernet Interface Overview	26
Configuring a Consistent Management IP Address	26
Configuring the MAC Address on the Management Ethernet Interface	28
Power over Ethernet (PoE) on ACX Series Routers	29
Understanding PoE on ACX Series Universal Metro Routers	30
ACX2000 PoE Specifications	30
PoE Classes and Power Ratings	31
PoE Options	31
Example: Configuring PoE on ACX2000 Routers	33
Example: Disabling a PoE Interface on ACX2000 Routers	38
Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers	40
Point-to-Point Protocol over Ethernet (PPPoE)	41
PPPoE Overview	41
PPPoE Interfaces	42
PPPoE Stages	43
Optional CHAP Authentication	44
Configuring PPPoE	45
Overview	46
Setting the Appropriate Encapsulation on the PPPoE Interface	47
Configuring PPPoE Encapsulation on an Ethernet Interface	48
Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface	48
Configuring the PPPoE Underlying Interface	48
Identifying the Access Concentrator	49
Configuring the PPPoE Automatic Reconnect Wait Timer	49
Configuring the PPPoE Service Name	49
Configuring the PPPoE Server Mode	50
Configuring the PPPoE Client Mode	50
Configuring the PPPoE Source and Destination Addresses	50



Deriving the PPPoE Source Address from a Specified Interface	50
Configuring the PPPoE IP Address by Negotiation	51
Configuring the Protocol MTU PPPoE	51
Example: Configuring a PPPoE Server Interface on an M120 or M320 Router	51
Disabling the Sending of PPPoE Keepalive Messages	52
Verifying a PPPoE Configuration	53
Tracing PPPoE Operations	53
Configuring the PPPoE Trace Log Filename	54
Configuring the Number and Size of PPPoE Log Files	55
Configuring Access to the PPPoE Log File	55
Configuring a Regular Expression for PPPoE Lines to Be Logged	55
Configuring the PPPoE Tracing Flags	56

## **Configuring Aggregated Ethernet Interfaces | 57**

Configuring Aggregated Ethernet Interfaces and LACP	57
Aggregated Ethernet Interfaces Overview	58
Platform Support for Aggregated Ethernet Interfaces	58
Enhanced LAG Support on MX Series Routers	60
Enhanced LAG Support on PTX Series Routers	61
Configuration Guidelines for Aggregated Ethernet Interfaces	61
Configuring an Aggregated Ethernet Interface	64
Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers	65
Support for Aggregated Ethernet Interfaces	66
Understanding Aggregated Ethernet Interfaces on Junos OS Evolved	66
How to Configure Aggregated Ethernet Interfaces	66
Configuring LACP for Aggregated Ethernet Interfaces	67
Configuring the LACP Interval	70
Configuring LACP Link Protection	70
Configuring LACP Hold-Up Timer to Prevent Link Flapping on LAG Interfaces	74
Tracing LACP Operations	75
Sample Configuration for Configuring Aggregated Ethernet LACP on Tagged and Untagged Interfaces	75
Configuring Junos OS for Supporting Aggregated Devices	77
Configuring Virtual Links for Aggregated Devices	77
Configuring LACP Link Protection at the Chassis Level	78

Enabling LACP Link Protection	79
Configuring System Priority	80
Configuring the Maximum Links Limit	80
Configuring PPM on Junos Fusion	80
Configuring the Number of Aggregated Ethernet Interfaces on the Device	81
Configuring Aggregated Ethernet Link Speed	81
Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces	84
Example: Configuring Aggregated Ethernet Interfaces	84
Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles	86
Aggregated Ethernet Bundle with Mixed Rates and Mixed Modes on T Series Routers	86
Aggregated Ethernet Bundles with Mixed Rates on MX Series Routers and PTX Series Routers	90
Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles	94
Configuring Mixed Rates and Mixed Modes on an Aggregated Ethernet Bundle on T Series Routers	95
Configuring Mixed Rates on Aggregated Ethernet Bundles on MX Series Routers	97
Configuring Tagged Aggregated Ethernet Interfaces	98
Configuring Untagged Aggregated Ethernet Interfaces	99
Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links	100
Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links Overview	100
Example: Configuring Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links	101
Understanding Independent Micro BFD Sessions for LAG	113
Example: Configuring Independent Micro BFD Sessions for LAG	116
Deleting an Aggregated Ethernet Interface	127
Link Protection of Aggregated Ethernet Interfaces	133
Configuring Aggregated Ethernet Link Protection	133
Configuring Link Protection for Aggregated Ethernet Interfaces	134
Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces	134
Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link	135
Disabling Link Protection for Aggregated Ethernet Interfaces	135
Configuring Aggregated Ethernet Minimum Links	135
Example: Configuring Aggregated Ethernet Link Protection	137

## Scheduling on Aggregated Ethernet Interfaces | 138

Configuring Shared Scheduling on Aggregated Ethernet Interfaces | 138

Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection | 139

## Load Balancing on Aggregated Ethernet Interfaces | 140

Load Balancing and Ethernet Link Aggregation Overview | 141

Understanding Aggregated Ethernet Load Balancing | 141

Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data | 144

Guidelines for Configuring Stateful Load Balancing for Aggregated Ethernet Interfaces or LAG Bundles | 146

Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | 147

Configuring Adaptive Load Balancing | 148

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers | 149

Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers Overview | 149

Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers | 150

Configuring Symmetrical Load Balancing on Trio-Based MPCs | 153

Example Configurations | 155

Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers | 156

Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers | 159

Configuring Symmetrical Hashing for family multiservice on Both Routers | 159

Configuring Symmetrical Hashing for family inet on Both Routers | 160

Configuring Symmetrical Hashing for family inet and family multiservice on the Two Routers | 160

Example: Configuring Aggregated Ethernet Load Balancing | 161

Example: Configuring Aggregated Ethernet Load Balancing | 162

## Performance Monitoring on Aggregated Ethernet Interfaces | 179

ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on Aggregated Ethernet Interfaces Overview | 179

Guidelines for Configuring Performance Monitoring Functionalities on Aggregated Ethernet Interfaces | 181

## Periodic Packet Management | 183

Understanding Periodic Packet Management on MX Series Routers | 184

Configuring Periodic Packet Management on MX Series Routers | 184

Identifying Periodic Packet Management Mode | 185

Enabling Centralized Periodic Packet Management | 186

## Understanding Ethernet Link Aggregation on ACX Series Routers | 188

Load Balancing | 190

LACP Monitoring | 191

Link Protection | 192

Configuring Link Protection for Aggregated Ethernet Interfaces | 193

Disabling Link Protection for Aggregated Ethernet Interfaces | 193

Understanding the Algorithm Used to Hash LAG Bundle | 193

## 2

## Gigabit Ethernet Interfaces

### Configuring Gigabit Ethernet Interfaces | 199

#### Configuring 10-Gigabit Ethernet PICs | 199

10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 200

12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 204

24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 206

P2-10G-40G-QSFPP PIC Overview | 207

Understanding Dual Configuration on P2-10G-40G-QSFPP PIC | 208

Understanding Port Group | 209

Port Numbering on P2-10G-40G-QSFPP PIC When Port Groups Are Not Configured | 213

10-Gigabit Ethernet Mode | 216

40-Gigabit Ethernet Mode | 217

#### Configuring the P2-10G-40G-QSFPP PIC | 218

Configuring the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 218

Configuring the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 218

Configuring the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 219

Configuring the PIC at Port Group Level | 220

Configuring Framing Mode on P2-10G-40G-QSFPP PIC | 220

Example: Configuring the P2-10G-40G-QSFPP PIC | 222

Framing Overview | 225

Understanding WAN Framing	227
Configuring Ethernet Framing	227
Modes of Operation	228
Configuring Mixed-Rate Mode Operation	229
Configuring Line-Rate Mode on PICs Supporting Oversubscription	230
Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC	231
Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC	232
Gigabit Ethernet Notification of Link Down Alarm Overview	236
Notification of Link Down for Optics Options Overview	236
Configuring Gigabit Ethernet Notification of Link Down Alarm	236
Configuring Link Down Notification for Optics Options Alarm or Warning	237
Configuring 40-Gigabit Ethernet PICs	238
40-Gigabit Ethernet PIC Overview	239
Configuring 40-Gigabit Ethernet PICs	241
Configuring 100-Gigabit Ethernet MICs/PICs	243
100-Gigabit Ethernet Interfaces Overview	243
MX Series 100-Gigabit Ethernet Interfaces	243
PTX Series 100-Gigabit Ethernet Interfaces	244
T Series 100-Gigabit Ethernet Interfaces	245
MPC3E MIC Overview	246
100-Gigabit Ethernet Type 4 PIC with CFP Overview	247
Configuring 100-Gigabit Ethernet Type 4 PIC With CFP	251
Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP	256
100-Gigabit Ethernet Type 5 PIC with CFP Overview	259
100-Gigabit Ethernet Interfaces Interoperability	261
Interoperability of the MIC-3D-1X100GE-CFP MIC with PICs on Other Routers	261
Interoperability of the MPC4E-3D-2CGE-8XGE MPC with PICs on Other Routers	262
Interoperability of the P1-PTX-2-100GE-CFP PIC with PICs on Other Routers	262
Interoperability of the PD-1CE-CFP-FPC4 PIC with PICs or MICs on Other Routers	262
Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP	263
Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4	264
Configuring SA Multicast Bit Steering Mode on the 100-Gigabit Ethernet PIC PF-1CGE-CFP	264

Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | **266**

Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | **267**

Configuring SA Multicast Bit Steering Mode on 100-Gigabit Ethernet PIC P1-PTX-2-100GE-CFP | **267**

Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface | **269**

Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | **271**

Transporting Legacy Traffic over Packet Switched Networks | **272**

Smart SFP Transceivers for Transporting PDH Traffic over PSNs Overview | **272**

Smart SFP Transceivers for Transporting SDH Traffic over PSNs Overview | **273**

Benefits of Smart SFP Transceivers | **273**

Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH Traffic | **274**

Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy SDH Traffic | **278**

Configuring Layer 2 Overhead Attribute in Interface Statistics | **283**

Accounting of the Layer 2 Overhead Attribute in Interface Statistics | **283**

Guidelines for Configuring the Computation of Layer 2 Overhead in Interface Statistics | **285**

Configuring Layer 2 Overhead Accounting in Interface Statistics | **286**

Enabling the Accounting of Layer 2 Overhead in Interface Statistics at the PIC Level | **286**

Verifying the Accounting of Layer 2 Overhead in Interface Statistics | **287**

Configuring Gigabit Ethernet Policers | **291**

Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | **291**

Configuring Gigabit Ethernet Policers | **293**

Overview | **294**

Configuring a Policer | **294**

Specifying an Input Priority Map | **295**

Specifying an Output Priority Map | **295**

Applying a Policer | **296**

Configuring MAC Address Filtering	298
Example: Configuring Gigabit Ethernet Policers	298
Configuring Gigabit Ethernet Two-Color and Tricolor Policers	300
Overview	301
Configuring a Policer	302
Applying a Policer	303
Example: Configuring and Applying a Policer	303
Gigabit Ethernet Autonegotiation	305
Gigabit Ethernet Autonegotiation Overview	305
Configuring Gigabit Ethernet Autonegotiation	306
Configuring Gigabit Ethernet Autonegotiation with Remote Fault	306
Configuring Flow Control	306
Configuring Autonegotiation Speed on MX Series Routers	306
Displaying Autonegotiation Status	307
<b>Configuring Rate Selectability  </b>	<b>313</b>
Introduction to Rate Selectability	313
Understanding Rate Selectability	314
Rate Selectability on MPC7E-MRATE	315
Rate Selectability on MIC-MRATE	315
Rate Selectability on JNP10K-LC2101	315
Rate Selectability on MIC-MACSEC-20GE	317
Guidelines for Configuring Rate Selectability	318
Guidelines for Configuring Rate Selectability for MIC-MRATE MIC and MPC7E-MRATE MPC	318
Guidelines for Configuring Rate Selectability for JNP10K-LC2101	318
Guidelines for Configuring Rate Selectability for MPC11E	319
MX10003 MPC Rate-Selectability Overview	320
Invalid Port Configuration	325
MX204 Router Rate-Selectability Overview	325
User-Configurable Rate Selectability of MX204 Routers	329
Maximum number of 10/40/100GE ports Configurable at PIC and Port Mode	330
Port Configuration - PIC Level	331
MPC10E-15C-MRATE Rate-Selectability Overview	332
User-Configurable Rate Selectability of MPC10E-15C-MRATE	335

MPC10E-10C-MRATE Rate-Selectability Overview | 336

    User-Configurable Rate Selectability of MPC10E-10C-MRATE | 338

PTX10003 Router Rate-Selectability Overview | 340

MX2K-MPC11E Rate Selectability Overview | 342

Interface Naming Conventions for Rate Selectability | 345

    Interface Naming Conventions for MPC7E-MRATE | 345

    Interface Naming Conventions for MIC-MRATE | 347

    Interface Naming Conventions for MX10003 MPC | 350

    Interface Naming Conventions for JNP10K-LC2101 | 352

    Interface Naming Conventions for MPC10E-15C-MRATE MPC | 355

    Interface Naming Conventions for MPC10E-10C-MRATE MPC | 356

    Interface Naming Conventions for MIC-MACSEC-20GE | 358

    Understanding Interface Naming Conventions for MX2K-MPC11E MPC | 360

Preventing Oversubscription Using Active Physical Ports | 364

    Supported Active Physical Ports for Configuring Rate Selectability to Prevent  
    Oversubscription | 364

    Supported Active Physical Ports for Configuring Rate Selectability to Prevent Oversubscription  
    on MX10003 MPC | 367

    Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204  
    Router | 369

        Invalid Port Configuration | 370

        Configuring Active Ports on MX204 Router with Rate Selectability | 370

    Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on  
    MPC10E-15C-MRATE | 373

    Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on  
    MPC10E-10C-MRATE | 376

Configuring Rate Selectability | 379

    Configuring Port Speed | 380

    Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers | 381

    Configuring Rate Selectability on MIC-MRATE to Enable Different Port Speeds | 385

        Configuring Rate Selectability on MIC-MRATE at MIC Level | 385

        Configuring Rate Selectability on MIC-MRATE at Port Level | 387

    Configuring Rate Selectability on MPC7E (Multi-Rate) to Enable Different Port Speeds | 389

        Configuring Rate Selectability at PIC Level | 390

        Configuring Rate Selectability at Port Level | 391



## Configuring Rate Selectability on MX10003 MPC to Enable Different Port Speeds | 394

Configuring Rate Selectability on MX10003 MPC at MIC/PIC Level | 394

Configuring Rate Selectability on MX10003 MPC at Port Level | 396

## Configuring Rate Selectability on MX204 to Enable Different Port Speeds | 398

Configuring Rate Selectability on MX204 at PIC Level | 398

Configuring Rate Selectability on MX204 at Port Level | 400

## Configuring Rate Selectability on JNP10K-2101 MPC to Enable Different Port Speeds | 402

Configuring Rate Selectability on JNP10K-2101 MPC at PIC Level | 402

Configuring Rate Selectability on JNP10K-LC2101 MPC at Port Level | 404

## Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds | 406

Configuring Rate Selectability on MPC10E-15C-MRATE at PIC Level | 406

Configuring Rate Selectability on MPC10E-15C-MRATE at Port Level | 408

## Configuring Rate Selectability on MPC10E-10C-MRATE to Enable Different Port Speeds | 410

Configuring Rate Selectability on MPC10E-10C-MRATE at PIC Level | 410

Configuring Rate Selectability on MPC10E-10C-MRATE at Port Level | 412

## Configuring Rate Selectability on the MX2K-MPC11E to Enable Different Port Speeds | 414

Configuring Rate Selectability on the MX2K-MPC11E at the PIC Level | 414

Configuring Rate Selectability on the MX2K-MPC11E at the Port Level | 415

## Configure Rate Selectability on ACX5448-D and ACX5448-M Routers | 417

## Channelize Interfaces on ACX5448-D and ACX5448-M Routers | 419

## Configuring Optical Transport network | 423

### Understanding Optical Transport Network (OTN) | 423

10-Gigabit Ethernet OTN Options Configuration Overview | 424

100-Gigabit Ethernet OTN Options Configuration Overview | 424

### Understanding the P1-PTX-24-10G-W-SFPP PIC | 426

Interface Features | 426

Layer 2 and Layer 3 Features | 428

OTN Alarms and Defects | 429

TCA Alarms | 430

### Understanding the features of ACX6360 | 430

Interface Features | 431

Section | ?

OTN Alarms and Defects | 431

TCA Alarms	432
Understanding the P2-100GE-OTN PIC	433
Interface Features	434
Layer 2 and Layer 3 Features	435
OTN Alarms and Defects	436
TCA Alarms	437
Understanding the MIC3-100G-DWDM MIC	437
Interface Features	438
Layer 2 and Layer 3 Features	439
OTN Alarms and Defects	439
Understanding the PTX-5-100G-WDM PIC	441
Interface Features	441
Layer 2 and Layer 3 Features	443
OTN Alarms and Defects	443
Understanding the PTX10K-LC1104 Line Card	445
Software Features	445
OTN Alarms and Defects	446
Interface Mapping and Modulation format for ACX6360	447
Interface Mapping and Modulation format for PTX10K-LC1104 Line Card	450
Supported OTN and Optics Options	452
Supported OTN Options on PTX Series Routers	453
Supported OTN Options on MX Series Routers	462
Supported OTN Options on ACX6360 Routers	471
Supported OTN Options on ACX5448-D Routers	477
Supported OTN Options on PTX10008 and PTX10016 Series Routers	482
Supported Optics Options on ACX6360 and ACX5448-D Routers	489
Supported Optics Options on PTX10008 and PTX10016 Series Routers	493
Forward Error Correction (FEC) and Bit Error Rate (BER)	497
Understanding Pre-FEC BER Monitoring and BER Thresholds	498
Supported Forward Error Correction Modes on MX Series Routers	502
Supported Forward Error Correction Modes on PTX Series Routers	503
Supported Forward Error Correction Modes on ACX6360 Router	504

Dense Wavelength Division Multiplexing (DWDM) Interface Wavelength | 505

Ethernet DWDM Interface Wavelength Overview | 505

Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength | 505

Configuring OTN | 508

Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC | 508

Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512

Configuring OTN Interfaces on P2-100GE-OTN PIC | 516

Configuring OTN Interfaces on MIC3-100G-DWDM MIC | 521

Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 526

Configuring OTN Interface Options on PTX10K-LC1104 | 532

ODU Path Delay Measurement for Performance Monitoring | 537

Understanding ODU Path Delay Measurement | 537

Guidelines for Configuring Delay Measurement | 538

Enabling ODU Path Delay Measurement | 539

Disabling ODU Path Delay Measurement | 541

### 3

## Operation, Administration, and Management (OAM) for Ethernet Interfaces

Configuring Connectivity Fault Management | 545

Ethernet Operations, Administration, and Maintenance | 545

Introduction to OAM Connectivity Fault Management (CFM) | 546

Ethernet OAM Connectivity Fault Management | 547

IEEE 802.1ag OAM Connectivity Fault Management Overview | 548

OAM Connectivity Fault Management Overview | 548

Connectivity Fault Management Key Elements | 550

Best Practices for Configuring 802.1ag Ethernet OAM for VPLS | 552

Configuring Connectivity Fault Management (CFM) | 553

Creating a Maintenance Domain | 555

Configuring Maintenance Intermediate Points (MIPs) | 556

Configuring Maintenance Association Intermediate Points in ACX Series | 557

Configuring the Maintenance Domain Bridge Domain | 558

Configuring the Maintenance Domain MIP Half Function | 559

Configuring the Maintenance Association Intermediate Points with Bridge Domain | 559

Configuring the Maintenance Association Intermediate Points with Circuit  
Cross-Connect | 559

Configuring the Maintenance Association Intermediate Points with Bridge Domain when Maintenance Association End Point is Configured	560
Configuring the Maintenance Intermediate Points with Circuit Cross-Connect when Maintenance Association End Point is Configured	560
Creating a Maintenance Association	561
Continuity Check Protocol Parameters Overview	562
Configuring Continuity Check Protocol Parameters for Fault Detection	563
Configuring a MEP to Generate and Respond to CFM Protocol Messages	564
Configuring a Maintenance Association End Point (MEP)	565
Configuring a remote Maintenance Association End Point (MEP)	567
Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements	568
Configuring Service Protection for VPWS over MPLS Using the MEP Interface	570
Configuring Linktrace Protocol in CFM	576
Configuring Rate Limiting of Ethernet OAM Messages	576
Configuring Ethernet Local Management Interface	579
Ethernet Local Management Interface Overview	580
Configuring the Ethernet Local Management Interface	582
Example E-LMI Configuration	584
Configuring a CFM Action Profile to Specify CFM Actions for CFM Events	589
CFM Action Profile to Bring Down a Group of Logical Interfaces Overview	591
Benefits of Creating CFM Action Profile to Bring Down a Group of Logical Interfaces	592
Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces	592
Enabling Enhanced Connectivity Fault Management Mode	596
Configuring M120 and MX Series Routers for CCC Encapsulated Packets	597
IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets Overview	597
CFM Features Supported on Layer 2 VPN Circuits	598
Configuring CFM for CCC Encapsulated Packets	598
Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades	599
Configuring Unified ISSU for 802.1ag CFM	600
Junos OS Support for Performance Monitoring Compliant with Technical Specification MEF 36	604

Damping CFM performance Monitoring Traps and Notifications to Prevent Congestion of The NMS | **605**

Example: Configuring Ethernet CFM on Physical Interfaces | **606**

Example: Configuring Ethernet CFM on Bridge Connections | **609**

Example: Configuring Ethernet CFM over VPLS | **614**

CFM Monitoring between CE and PE Devices | **624**

Understanding CFM Monitoring between CE and PE Devices | **625**

Single Active Multi-homing Use Case using RDI bit | **626**

Active/Active Multihoming Use case using RDI bit | **626**

Configuring Port Status TLV and Interface Status TLV | **626**

TLVs Overview | **627**

Various TLVs for CFM PDUs | **627**

Support for Additional Optional TLVs | **629**

MAC Status Defects | **636**

Configuring Remote MEP Action Profile Support | **638**

Monitoring a Remote MEP Action Profile | **639**

Configuring Chassis ID TLV | **640**

Configuring MAC Flush Message Processing in CET Mode | **642**

Configuring a Connection Protection TLV Action Profile | **644**

Example: Configuring an Action Profile Based on Connection Protection TLVs | **645**

Configuring Continuity Check Messages | **648**

Configuring Faster Protection Switching for Point-to-Point Network Topologies | **649**

Configuring Faster Convergence for Dual-Homed Multipoint-to-Multipoint Network Topologies | **651**

Configuring a Primary VLAN ID for Increased Flexibility | **652**

Configuring a Remote Maintenance Association to Accept a Different ID | **653**

**Configuring Link Fault Management | 655**

Introduction to OAM Link Fault Management (LFM) | **655**

IEEE 802.3ah OAM Link-Fault Management Overview | **655**

Understanding Ethernet OAM Link Fault Management for ACX Series Routers | **656**

Configuring Ethernet 802.3ah OAM | **658**

Configuring Link Fault Management | **660**

Configuring Link Discovery | **661**

Configuring the OAM PDU Interval | **662**

Configuring the OAM PDU Threshold	662
Configuring Threshold Values for Local Fault Events on an Interface	662
Disabling the Sending of Link Event TLVs	663
Example: Configuring IEEE 802.3ah OAM Support on an Interface	663
Example: Configuring IEEE 802.3ah OAM Support for an Interface on ACX Series	664
Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge	667
Example: Configuring Ethernet LFM for CCC	669
Example: Configuring Ethernet LFM for Aggregated Ethernet	670
Configuring an OAM Action Profile	673
Specifying the Actions to Be Taken for Link-Fault Management Events	674
Monitoring the Loss of Link Adjacency	675
Monitoring Protocol Status	676
Configuring Threshold Values for Fault Events in an Action Profile	677
Applying an Action Profile	677

## Remote Fault Detection for Link Fault Management | 678

Detecting Remote Faults	678
Enabling Dying Gasp Functionality	679

## Remote Loopback for Link Fault Management | 680

Setting a Remote Interface into Loopback Mode	680
Enabling Remote Loopback Support on the Local Interface	681
Enabling Nonstop Routing for Ethernet Link Fault Management on Backup Routers	682
Example: Configuring Ethernet LFM with Loopback Support	686

## Configuring ITU-TY.1731 Ethernet Service OAM | 689

### ITU-T Y.1731 Ethernet Service OAM Overview | 689

Ethernet Frame Delay Measurements Overview	690
ITU-T Y.1731 Frame Delay Measurement Feature	690
One-Way Ethernet Frame Delay Measurement	692
Two-Way Ethernet Frame Delay Measurement	693
Choosing Between One-Way and Two-Way ETH-DM	695
Restrictions for Ethernet Frame Delay Measurement	695
Ethernet Frame Loss Measurement Overview	696
Service-Level Agreement Measurement	698
On-Demand Mode for SLA Measurement	699

Proactive Mode for SLA Measurement	700
Ethernet Delay Measurements and Loss Measurement by Proactive Mode	701
Ethernet Failure Notification Protocol Overview	701
Ethernet Synthetic Loss Measurement Overview	702
Scenarios for Configuration of ETH-SLM	703
Upstream MEP in MPLS Tunnels	704
Downstream MEP in Ethernet Networks	704
Format of ETH-SLM Messages	705
SLM PDU Format	705
SLR PDU Format	705
Data Iterator TLV Format	706
Transmission of ETH-SLM Messages	706
Initiation and Transmission of SLM Requests	707
Reception of SLMs and Transmission of SLRs	707
Reception of SLRs	708
Computation of Frame Loss	708
Configuring Ethernet Frame Delay Measurement Sessions	709
Guidelines for Configuring Routers to Support an ETH-DM Session	710
Configuration Requirements for ETH-DM	710
Configuration Options for ETH-DM	711
Guidelines for Starting an ETH-DM Session	711
ETH-DM Session Prerequisites	711
ETH-DM Session Parameters	712
Restrictions for an ETH-DM Session	713
Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts	713
ETH-DM Statistics	714
ETH-DM Statistics Retrieval	716
ETH-DM Frame Counts	716
ETH-DM Frame Count Retrieval	717
Configuring Routers to Support an ETH-DM Session	718
Configuring MEP Interfaces	718
Ensuring That Distributed ppm Is Not Disabled	719
Enabling the Hardware-Assisted Timestamping Option	722

Configuring the Server-Side Processing Option	722
Triggering an Ethernet Frame Delay Measurements Session	723
Starting an ETH-DM Session	725
Using the monitor ethernet delay-measurement Command	725
Starting a One-Way ETH-DM Session	726
Starting a Two-Way ETH-DM Session	727
Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces	728
Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces	734
Managing Continuity Measurement Statistics	740
Displaying Continuity Measurement Statistics	740
Clearing Continuity Measurement Statistics	741
Viewing Ethernet Frame Delay Measurements Statistics	741
Managing ETH-DM Statistics and ETH-DM Frame Counts	743
Displaying ETH-DM Statistics Only	743
Displaying ETH-DM Statistics and Frame Counts	744
Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity	744
Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level	745
Clearing ETH-DM Statistics and Frame Counts	746
Configuring Ethernet Frame Loss Measurement	747
Configuring Statistical Frame Loss Measurement for VPLS Connections	747
Managing ETH-LM Statistics	748
Displaying ETH-LM Statistics	749
Clearing ETH-LM Statistics	750
Example: Measuring Ethernet Frame Loss for Single-Tagged LMM/LMR PDUs	750
Example: Measuring Ethernet Frame Loss for Dual-Tagged LMM/LMR PDUs	766
Configuring an Iterator Profile	781
Configuring an Iterator Profile	782
Verifying the Configuration of an Iterator Profile	785
Displaying the Configuration of an Iterator Profile for Two-way Delay Measurement	785
Displaying the Configuration of an Iterator Profile for Loss Measurement	786
Displaying the Configuration of a Remote MEP with an Iterator Profile	787



Disabling an Iterator Profile	787
Managing Iterator Statistics	788
Displaying Iterator Statistics	788
Clearing Iterator Statistics	794
Configuring a Remote MEP with an Iterator Profile	794
Configuring Ethernet Synthetic Loss Measurements	796
Guidelines for Configuring ETH-SLM	797
Starting a Proactive ETH-SLM Session	798
Configuring MEP Interfaces	798
Configuring an Iterator Profile for ETH-SLM	799
Associating the Iterator Profile with MEPs for ETH-SLM	801
Starting an On-Demand ETH-SLM Session	803
Managing ETH-SLM Statistics and ETH-SLM Frame Counts	804
Displaying ETH-SLM Statistics Only	804
Displaying ETH-SLM Statistics and Frame Counts	805
Displaying ETH-SLM Frame Counts for MEPs by Enclosing CFM Entity	806
Displaying ETH-SLM Frame Counts for MEPs by Interface or Domain Level	807
Clearing ETH-SLM Statistics and Frame Counts	808
Clearing Iterator Statistics	808
Troubleshooting Failures with ETH-SLM	809
Ethernet Alarm Indication	810
Ethernet Alarm Indication Signal (ETH-AIS) Function Overview	810
Understanding ETH-AIS in a Maintenance Domain	811
Fault Detection in a Maintenance Domain	811
Terms Defined	813
Ethernet Alarm Indication Signal Overview	815
Configuring ETH-AIS on a CFM MEP	816
Configuring an Action Profile	817
Configuring an Action to Be Taken When an AIS Alarm Is Detected	818
Attaching the Action Profile to a CFM MEP	819
Configuring Alarm Indication Signal on ACX Series Routers	821
Inline Transmission Mode	824
Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling	825
Enabling Inline Transmission of Link Fault Management Keepalives for Maximum Scaling	826

Enabling Inline Mode Of Performance Monitoring To Achieve Maximum Scaling | 830

Supported Inline CCM and Inline PM Scaling Values | 833

Enabling the Hardware-Assisted Timestamping Option | 835

## Troubleshooting Information

### Monitoring and Troubleshooting Ethernet Interfaces | 839

Passive Monitoring on Ethernet Interfaces Overview | 839

Enabling Passive Monitoring on Ethernet Interfaces | 841

Link Degrade Monitoring Overview | 843

Supported Platforms | 844

Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 846

Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 846

Monitor Fast Ethernet and Gigabit Ethernet Interfaces | 847

Display the Status of Fast Ethernet Interfaces | 847

Display the Status of Gigabit Ethernet Interfaces | 849

Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface | 850

Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface | 852

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 855

Fiber-Optic Ethernet Interface Specifications | 857

Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 859

Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 859

Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface | 860

Create a Loopback | 861

Create a Physical Loopback for a Fiber-Optic Interface | 861

Create a Loopback Plug for an RJ-45 Ethernet Interface | 862

Configure a Local Loopback | 863

Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 864

Configure a Static Address Resolution Protocol Table Entry | 868

Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 873

Ping the Fast Ethernet or Gigabit Ethernet Interface | 874

Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 875

Diagnose a Suspected Circuit Problem | 877

- Performing Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 878
  - Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 878
  - Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | 879
  - Create a Loopback | 880
    - Create a Physical Loopback for a Fiber-Optic Interface | 880
    - Configure a Local Loopback | 881
  - Verify That the Gigabit Ethernet Interface Is Up | 882
  - Configure a Static Address Resolution Protocol Table Entry | 887
  - Clear Gigabit Ethernet Interface Statistics | 889
  - Ping the Gigabit Ethernet Interface | 889
  - Check for Gigabit Ethernet Interface Error Statistics | 890
  - Diagnose a Suspected Circuit Problem | 894
- Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 895
  - Configuring Loopback Testing | 895
  - Configuring BERT Testing | 897
  - Starting and Stopping a BERT Test | 901
- Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters | 902
  - Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters | 902
  - Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm | 903
  - Fast Ethernet and Gigabit Ethernet Counters | 905
- Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State | 908
- Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | 910
  - Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test | 910
  - Guidelines to perform Bidirectional Diagnostics using Remote Loopback | 914
  - Clearing the Interface Statistics | 917

## Configuration Statements and Operational Commands

### Configuration Statements (OTN) | 923

- alarm (optics-options) | 925
- backward-frr-enable | 926
- ber-threshold-clear | 928
- ber-threshold-signal-degrade | 931
- bypass | 934
- bytes (otn-options) | 935

cfp-to-et | 936  
fec | 937  
fec (gigether) | 940  
fixed-stuff-bytes | 941  
high-polarization | 942  
interval | 943  
is-ma | 945  
laser-enable | 946  
line-loopback | 947  
local-loopback | 948  
monitor-end-point | 949  
no-odu-backward-frr-enable | 950  
no-odu-signal-degrade-monitor-enable | 951  
number-of-frames | 952  
oc192 | 953  
odu-delay-management | 954  
odu-backward-frr-enable | 955  
odu-signal-degrade | 956  
odu-signal-degrade-monitor-enable | 958  
odu-ttim-action-enable | 959  
otu-ttim-action-enable | 960  
otu4 | 961  
pass-through | 962  
prbs | 963  
preemptive-fast-reroute | 964  
rate | 965  
remote-loop-enable | 966  
signal-degrade | 968  
signal-degrade-monitor-enable | 969  
start-measurement | 971  
tca | 972  
transport-monitoring | 975  
trigger | 976  
tti | 981

tx-power | 982

warning | 983

wavelength | 984

## **Configuration Statements (OAM-CFM) | 989**

802.3ad | 992

action-profile (Applying to CFM) | 993

action-profile (Defining for CFM) | 994

action-profile (MEP) | 996

ais-trigger-condition | 997

all-defects | 998

auto-discovery | 999

avg-fd-twoway-threshold | 1000

avg-ifdv-twoway-threshold | 1002

avg-flr-forward-threshold | 1004

avg-flr-backward-threshold | 1006

calculation-weight | 1008

clear-action (CFM) | 1009

connection-protection-tlv | 1010

continuity-check | 1011

convey-loss-threshold | 1012

cross-connect-ccm | 1013

cycle-time | 1014

data-tlv-size | 1015

default-actions | 1016

delay | 1017

delegate-server-processing | 1018

delay-variation | 1019

detect-loc | 1020

direction | 1022

enhanced-cfm-mode | 1023

erroneous-ccm | 1024

event (CFM) | 1025

flap-trap-monitor | 1026

hardware-assisted-timestamping | 1028

hardware-assisted-keepalives | 1029

hold-interval (OAM) | 1031

instance | 1032

interface (IEEE 802.1ag OAM Connectivity-Fault Management) | 1033

interface (OAM Link-Fault Management) | 1034

interface-down | 1035

interface-status-tlv | 1036

interface-status-send-rdi | 1037

interval | 1038

interval (CFM MEP) | 1040

iteration-count | 1041

iteration-period | 1042

level | 1043

level (CFM MEP) | 1044

linktrace | 1045

log-and-generate-ais | 1046

loss-threshold | 1047

lowest-priority-defect | 1048

maintenance-association | 1049

maintenance-domain | 1051

measurement-interval | 1053

measurement-type | 1054

mep | 1055

mip-half-function | 1057

name-format | 1058

no-aggregate-delegate-processing | 1059

path-database-size | 1060

performance-monitoring | 1061

policer (CFM Global) | 1063

policer (CFM Session) | 1064

port-status-tlv | 1065

priority (Protocols OAM) | 1066

priority (CFM MEP) | 1067

priority (OAM Connectivity-Fault Management) | 1068

protocol (Server MEP) | 1069

protect-maintenance-association (OAM) | 1070

rdi | 1071

receive-ais | 1072

remote-maintenance-association (OAM) | 1073

remote-mep | 1074

sendid-tlv | 1075

short-name-format | 1076

sla-iterator-profile | 1077

sla-iterator-profiles | 1079

## **Configuration Statements | 1081**

accept-source-mac | 1088

access-concentrator | 1090

account-layer2-overhead (PIC Level) | 1091

action (OAM) | 1092

action-profile | 1093

adaptive | 1095

address | 1097

adjacency-loss | 1101

age | 1102

agent-specifier | 1103

aggregate (Gigabit Ethernet CoS Policer) | 1105

aggregated-devices | 1106

aggregated-ether-options | 1108

alarms | 1110

allow-remote-loopback | 1111

apply-action-profile | 1112

asynchronous-notification | 1113

auto-negotiation | 1114

auto-reconnect | 1116

bandwidth-limit (Policer for Gigabit Ethernet Interfaces) | 1117

bert-algorithm | 1118

bert-error-rate | **1121**

bert-period | **1123**

bridge-domain | **1125**

burst-size-limit (Policer for Gigabit Ethernet Interfaces) | **1126**

centralized | **1127**

cfp-to-et | **1128**

classifier | **1129**

client | **1130**

community-vlans (MX Series) | **1131**

connectivity-fault-management | **1132**

delay (PPPoE Service Name Tables) | **1136**

destination (IPCP) | **1137**

device-count | **1138**

disable (Link Protection) | **1139**

distribution-list | **1140**

drop (PPPoE Service Name Tables) | **1141**

dynamic-profile (PPPoE Service Name Tables) | **1142**

egress-policer-overhead | **1143**

encapsulation (Logical Interface) | **1145**

encapsulation | **1150**

ether-options | **1157**

ethernet (Chassis) | **1159**

ethernet (Protocols OAM) | **1160**

ethernet-policer-profile | **1168**

evcs | **1170**

evc-protocol cfm | **1171**

event (LFM) | **1173**

event-thresholds | **1175**

family | **1176**

fastether-options | **1182**

flow-control | **1184**

fnp | **1186**

force-up | **1187**

forwarding-class (Gigabit Ethernet IQ Classifier) | **1188**



forwarding-mode (100-Gigabit Ethernet) | **1189**

forwarding-mode (PTX Series Packet Transport Routers) | **1190**

frame-error | **1191**

frame-period | **1192**

frame-period-summary | **1193**

framing (10-Gigabit Ethernet Interfaces) | **1194**

gether-options | **1196**

hold-time up | **1198**

ieee802.1p | **1199**

ignore-l3-incompletes | **1200**

ingress-policer-overhead | **1201**

ingress-rate-limit | **1203**

inline | **1204**

input-policer | **1206**

input-priority-map | **1207**

input-three-color | **1208**

input-vlan-map (Aggregated Ethernet) | **1209**

interface | **1210**

interface (OAM Link-Fault Management) | **1212**

interface-group | **1214**

interface-group-down | **1215**

interface-none | **1216**

isolated-vlan (MX Series) | **1217**

lACP (802.3ad) | **1218**

lACP (Aggregated Ethernet) | **1219**

lACP (Protocols) | **1222**

layer2-policer | **1224**

link-adjacency-loss | **1225**

link-discovery | **1226**

link-degrade-monitor | **1227**

link-down | **1228**

link-event-rate | **1229**

link-fault-management | **1230**

link-mode | **1232**

link-protection | **1234**

link-protection (non-LACP) | **1236**

link-speed (Aggregated Ethernet) | **1238**

link-speed (Aggregated SONET/SDH) | **1240**

lmi (Ethernet OAM) | **1241**

load-balance | **1243**

load-balance-stateful (Aggregated Ethernet Interfaces) | **1245**

load-type (Aggregated Ethernet Interfaces) | **1246**

local-bias (ae load-balance) | **1247**

logical-tunnel-options | **1249**

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | **1251**

loopback (Local and Remote) | **1253**

loopback-tracking | **1254**

loss-priority | **1255**

mac | **1256**

mac-address (Accept Source Mac) | **1257**

mac-learn-enable | **1258**

mac-validate | **1259**

master-only | **1260**

max-sessions (PPPoE Service Name Tables) | **1261**

max-sessions-vs-a-ignore (Static and Dynamic Subscribers) | **1262**

maximum-links | **1264**

mc-ae | **1266**

minimum-bandwidth (aggregated Ethernet) | **1272**

minimum-links | **1273**

mixed-rate-mode | **1274**

mtu | **1275**

mru | **1279**

multicast-statistics | **1280**

negotiate-address | **1281**

negotiation-options | **1282**

no-adaptive | **1283**

no-allow-link-events | **1284**

no-auto-mdix | **1285**

no-keepalives | **1286**  
no-pre-classifier | **1287**  
no-send-pads-ac-info | **1288**  
no-send-pads-error | **1289**  
non-revertive (Interfaces) | **1290**  
number-of-sub-ports | **1291**  
oam | **1293**  
optics-options | **1297**  
otn-options | **1299**  
output-policer | **1301**  
output-priority-map | **1302**  
output-three-color | **1303**  
output-vlan-map (Aggregated Ethernet) | **1304**  
pado-advertise | **1305**  
passive-monitor-mode | **1306**  
pdu-interval | **1307**  
pdu-threshold | **1308**  
per-flow (Aggregated Ethernet Interfaces) | **1309**  
periodic | **1310**  
policer (CFM Firewall) | **1312**  
policer (CoS) | **1313**  
policer (MAC) | **1315**  
port-priority | **1317**  
pp0 (Dynamic PPPoE) | **1318**  
ppm (Ethernet Switching) | **1320**  
pppoe-options | **1321**  
pppoe-underlying-options (Static and Dynamic Subscribers) | **1323**  
preferred-source-address | **1324**  
premium (Output Priority Map) | **1326**  
premium (Policer) | **1327**  
protocol-down | **1328**  
premium (Output Priority Map) | **1329**  
premium (Policer) | **1330**  
proxy | **1331**

rebalance (Aggregated Ethernet Interfaces) | 1332

receive-options-packets | 1333

receive-ttl-exceeded | 1334

recovery | 1335

remote-loopback | 1337

restore-interval | 1338

revertive | 1339

routing-instance | 1340

routing-instance (PPPoE Service Name Tables) | 1341

rx-enable | 1342

rx-max-duration | 1343

sa-multicast (100-Gigabit Ethernet) | 1344

sa-multicast (PTX Series Packet Transport Routers) | 1346

send-critical-event | 1347

server | 1348

service (PPPoE) | 1349

service-name | 1351

service-name-table | 1352

service-name-tables | 1354

session-expiry (MX Series in Enhanced LAN Mode) | 1355

source-address-filter | 1356

source-filtering | 1358

speed (Ethernet) | 1359

speed (MX Series DPC) | 1368

static-interface | 1369

switch-options | 1370

switch-port | 1371

symbol-period | 1372

sync-reset | 1373

syslog (OAM Action) | 1374

system-id | 1375

system-priority | 1376

targeted-options (Grouping Subscribers by Bandwidth Usage) | 1377

targeted-options (Manual Targeting) | 1379

targeted-distribution | **1381**  
targeted-options | **1382**  
terminate (PPPoE Service Name Tables) | **1383**  
thresholds | **1384**  
traceoptions (LLDP) | **1386**  
traceoptions (Individual Interfaces) | **1389**  
traceoptions (LACP) | **1398**  
traceoptions (PPPoE) | **1400**  
tx-duration | **1403**  
tx-enable | **1404**  
underlying-interface | **1405**  
unit | **1407**  
unnumbered-address (Dynamic Profiles) | **1418**  
unnumbered-address (PPP) | **1420**  
Unused | **1421**  
virtual-control-channel | **1422**  
virtual-switch | **1423**  
vlan-rule (100-Gigabit Ethernet Type 4 PIC with CFP) | **1424**  
vlan-steering (100-Gigabit Ethernet Type 4 PIC with CFP) | **1425**

## **Operational Commands | 1427**

clear interfaces interface-set statistics | **1430**  
clear interfaces interval | **1431**  
clear interfaces aeX forwarding-options load-balance state | **1433**  
clear interfaces aggregate forwarding-options load-balance state | **1434**  
clear interfaces transport pm | **1435**  
clear oam ethernet connectivity-fault-management continuity-measurement | **1437**  
clear oam ethernet connectivity-fault-management linktrace path-database | **1439**  
clear oam ethernet connectivity-fault-management loss-statistics | **1440**  
clear oam ethernet connectivity-fault-management policer | **1442**  
clear oam ethernet connectivity-fault-management statistics | **1444**  
clear oam ethernet connectivity-fault-management synthetic-loss-measurement | **1446**  
clear oam ethernet link-fault-management state | **1448**  
clear oam ethernet link-fault-management statistics | **1449**

clear protection-group ethernet-ring statistics | **1450**

monitor ethernet delay-measurement | **1451**

monitor ethernet loss-measurement | **1457**

monitor ethernet synthetic-loss-measurement | **1463**

monitor ethernet synthetic-loss-measurement | **1468**

prbs-test-start | **1473**

prbs-test-stop | **1475**

request interface link-degrade-recover | **1476**

request interface mc-ae switchover (Multichassis Link Aggregation) | **1480**

request interface (revert | switchover) (Aggregated Ethernet Link Protection) | **1482**

request lacp link-switchover | **1484**

show chassis hardware | **1485**

show chassis pic | **1504**

show ethernet-switching redundancy-groups | **1540**

show interfaces (Adaptive Services) | **1545**

show interfaces (Aggregated Ethernet) | **1553**

show interfaces demux0 (Demux Interfaces) | **1568**

show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port) | **1580**

show interfaces (far-end-interval) | **1631**

show interfaces (Fast Ethernet) | **1633**

show interfaces | **1656**

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) | **1763**

show interfaces (PPPoE) | **1786**

show interfaces interface-set (Ethernet Interface Set) | **1799**

show interfaces interface-set queue | **1805**

show interfaces interval | **1816**

show interfaces irb | **1820**

show interfaces mac-database | **1828**

show interfaces mc-ae | **1835**

show interfaces prbs-stats | **1842**

show interfaces transport pm | **1844**

show l2-learning instance | **1854**

show l2-learning redundancy-groups | **1856**

show lacp interfaces | **1862**

show oam ethernet connectivity-fault-management delay-statistics | **1869**

show oam ethernet connectivity-fault-management forwarding-state | **1874**

show oam ethernet connectivity-fault-management interfaces | **1879**

show oam ethernet connectivity-fault-management linktrace path-database | **1893**

show oam ethernet connectivity-fault-management loss-statistics | **1896**

show oam ethernet connectivity-fault-management mep-database | **1901**

show oam ethernet connectivity-fault-management mep-statistics | **1916**

show oam ethernet connectivity-fault-management path-database | **1932**

show oam ethernet connectivity-fault-management policer | **1934**

show oam ethernet connectivity-fault-management sla-iterator-statistics | **1938**

show oam ethernet connectivity-fault-management synthetic-loss-statistics | **1951**

show oam ethernet evc | **1954**

show oam ethernet fnp interface | **1956**

show oam ethernet fnp messages | **1958**

show oam ethernet fnp status | **1960**

show oam ethernet link-fault-management | **1962**

show oam ethernet lmi | **1972**

show oam ethernet lmi statistics | **1974**

show pppoe interfaces | **1976**

show pppoe service-name-tables | **1981**

show pppoe sessions | **1985**

show pppoe statistics | **1987**

show pppoe underlying-interfaces | **1990**

show pppoe version | **1999**

show protection-group ethernet-ring aps | **2001**

show protection-group ethernet-ring configuration | **2006**

show protection-group ethernet-ring data-channel | **2015**

show protection-group ethernet-ring flush-info | **2018**

show protection-group ethernet-ring interface | **2020**

show protection-group ethernet-ring node-state | **2025**

show protection-group ethernet-ring statistics | **2031**

show protection-group ethernet-ring vlan | **2038**

traceroute ethernet | **2044**

# About the Documentation

## IN THIS SECTION

- Documentation and Release Notes | xxxvii
- Using the Examples in This Manual | xxxvii
- Documentation Conventions | xxxix
- Documentation Feedback | xlii
- Requesting Technical Support | xlii

Use this guide to configure, monitor, and troubleshoot the various supported Ethernet Interfaces, including aggregated Ethernet Interfaces on Juniper Networks routers. You can also configure Ethernet Operations, Administration, and Maintenance (OAM) on Juniper Networks routers to determine how the network of Ethernet links is functioning.

## Documentation and Release Notes

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

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

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

## Using the Examples in This Manual

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



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

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

## Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

## Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see [CLI Explorer](#).

## Documentation Conventions

[Table 1 on page xl](#) defines notice icons used in this guide.

Table 1: Notice Icons

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

Table 2 on page xl defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>

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

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>
<b>Text like this</b>	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit <b>protocols ospf area area-id</b>] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast   multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ <i>community-ids</i> ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

## GUI Conventions


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

Convention	Description	Examples
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

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



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

## Self-Help Online Tools and Resources

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

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

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

## Creating a Service Request with JTAC

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

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

# 1

PART

## Ethernet Interfaces

---

Configuring Ethernet Interfaces | 3

Configuring Aggregated Ethernet Interfaces | 57

---





# Configuring Ethernet Interfaces

## IN THIS CHAPTER

- [Ethernet Interfaces Overview | 3](#)
- [Initial Configuration of Ethernet Interfaces | 4](#)
- [Flow Control for Ethernet Interfaces | 17](#)
- [MAC Address Filtering and Accounting on Ethernet Interfaces | 20](#)
- [Management Ethernet Interfaces | 25](#)
- [Power over Ethernet \(PoE\) on ACX Series Routers | 29](#)
- [Point-to-Point Protocol over Ethernet \(PPPoE\) | 41](#)

## 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 | 21](#)

[Configuring Flow Control | 19](#)

[Configuring the Interface Speed on Ethernet Interfaces | 7](#)

## Initial Configuration of Ethernet Interfaces

### IN THIS SECTION

- [Configuring Ethernet Physical Interface Properties | 5](#)
- [Configuring the Interface Speed on Ethernet Interfaces | 7](#)
- [Configuring the Ingress Rate Limit | 8](#)
- [Configuring the Link Characteristics on Ethernet Interfaces | 9](#)
- [Configuring Multicast Statistics Collection on Ethernet Interfaces | 10](#)
- [MAC Address Validation on Static Ethernet Interfaces Overview | 11](#)
- [Configuring MAC Address Validation on Static Ethernet Interfaces | 12](#)
- [Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router | 13](#)
- [Example: Configuring Fast Ethernet Interfaces | 15](#)
- [Example: Configuring Gigabit Ethernet Interfaces | 16](#)

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 **gigether-options** statement at the **[edit interfaces ge-fpc/pic/port]** or **[edit interfaces xe-fpc/pic/port]** hierarchy level:

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

- 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 {
```

- 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;
```

- 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;
```

- 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;
```

- 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 | 16](#)

## 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 RouterMX 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](#) | [1359](#)

## 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](#) | [1203](#)

## 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](#) | [1232](#)

## 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:

1. 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 | 1280](#)

[Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces | 84](#)

## 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

| [Disabling MAC Address Learning of Neighbors Through ARP or Neighbor Discovery for IPv4 and IPv6 Neighbors](#)

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

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

SEE ALSO

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 | 25](#)

*Configuring Ethernet Interfaces*

### Example: Configuring Gigabit Ethernet Interfaces

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

```

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

```

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

Each of the two interfaces on the PIC is named:

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

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

SEE ALSO

<a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces</a>
<a href="#">Display the Status of Gigabit Ethernet Interfaces   849</a>
<a href="#">Monitoring Fast Ethernet and Gigabit Ethernet Interfaces   846</a>

Release History Table

Release	Description
<a href="#">14.2</a>	Starting with Junos OS Release 14.2 the <b>auto-10m-100m</b> option allows the fixed tri-speed port to auto negotiate with ports limited by <b>100m</b> or <b>10m</b> 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.
<a href="#">11.4</a>	Starting with Junos OS Release 11.4, half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces. When you include the <b>speed</b> statement, you must include the <b>link-mode full-duplex</b> statement at the same hierarchy level.

RELATED DOCUMENTATION

<a href="#">MAC Address Filtering and Accounting on Ethernet Interfaces   20</a>
<a href="#">Management Ethernet Interfaces   25</a>

Flow Control for Ethernet Interfaces

IN THIS SECTION

- [Understanding Flow Control | 18](#)
- [Configuring Flow Control | 19](#)

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 | 18](#)

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](#) | [1184](#)

### **Configuring Flow Control**

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

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

```
no-flow-control;
```

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

```
flow-control;
```

You can include these statements at the following hierarchy levels:

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

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

#### SEE ALSO

[flow-control | 1184](#)

[Ethernet Interfaces Overview | 3](#)

#### 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 | 21](#)
- [Configuring MAC Address Filtering on PTX Series Packet Transport Routers | 23](#)
- [Configuring MAC Address Accounting | 24](#)

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 | 21](#)

### *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” on page 293](#).

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 [“Configuring LACP for Aggregated Ethernet Interfaces” on page 67](#).

**NOTE:** On untagged Gigabit Ethernet interfaces, you should not configure the **source-address-filter** statement at the **[edit interfaces ge-fpc/pic/port together-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-fpc/pic/port together-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 together-options]** hierarchy level and the **accept-source-mac** statement at the **[edit interfaces ge-fpc/pic/port together-options unit logical-unit-number]** hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

**NOTE:** The **source-address-filter** statement is not supported on MX Series routers with MPC4E (model numbers: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE); instead, include the **accept-source-mac** statement. For more information, see [“Configuring Gigabit Ethernet Policers” on page 293](#).

SEE ALSO

[source-address-filter | 1356](#)

[Configuring LACP for Aggregated Ethernet Interfaces | 67](#)

[Ethernet Interfaces Overview | 3](#)

## 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 together-options]** hierarchy level:

```
[edit interfaces]
et-x/y/z {
  together-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 {
  together-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](#) | 1828

### 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 | 291](#)

[Configuring Gigabit Ethernet Policers | 293](#)

[Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 300](#)

[Configuring a Policer Overhead](#)

#### RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Interfaces and LACP | 57](#)

[Configuring Gigabit Ethernet Policers | 291](#)

## Management Ethernet Interfaces

#### IN THIS SECTION

- [Management Ethernet Interface Overview | 26](#)
- [Configuring a Consistent Management IP Address | 26](#)
- [Configuring the MAC Address on the Management Ethernet Interface | 28](#)

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 | 3](#)

[Initial Configuration of Ethernet Interfaces | 4](#)

## 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 master Routing Engine, you must know which Routing Engine is active and use the appropriate IP address.

Optionally, for consistent access to the master 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 master Routing Engine. During switchover, the address moves to the new master 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 master Routing Engine. The address remains consistent regardless of which Routing Engine is active. IP address **10.17.40.132** is assigned to **fxp0** on **re0**, and address **10.17.40.133** is assigned to **fxp0** on **re1**.

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

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

SEE ALSO

[Ethernet Interfaces Overview | 3](#)

[Initial Configuration of Ethernet Interfaces | 4](#)

## 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 | 3](#)

[Initial Configuration of Ethernet Interfaces | 4](#)

[mac | 1256](#)

#### RELATED DOCUMENTATION

[Ethernet Interfaces Overview | 3](#)

[Initial Configuration of Ethernet Interfaces | 4](#)

## Power over Ethernet (PoE) on ACX Series Routers

#### IN THIS SECTION

- [Understanding PoE on ACX Series Universal Metro Routers | 30](#)
- [Example: Configuring PoE on ACX2000 Routers | 33](#)
- [Example: Disabling a PoE Interface on ACX2000 Routers | 38](#)
- [Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers | 40](#)

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 | 31](#)
- [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 3 on page 30](#) lists the PoE specifications for the ACX2000 routers.

Table 3: PoE Specifications for the ACX2000 Routers

Specifications	For ACX2000 Universal Metro Routers
Supported standards	<ul style="list-style-type: none"><li>● IEEE 802.3 AF</li><li>● IEEE 802.3 AT (PoE+)</li><li>● Legacy (pre-standards)</li></ul>
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

Table 3: PoE Specifications for the ACX2000 Routers (*continued*)

Specifications	For ACX2000 Universal Metro Routers
Default per port power limit	32 W
Maximum per port power limit	65 W
Power management modes	<ul style="list-style-type: none"> <li>• <b>class</b>—Power allocated for each interface can be configured.</li> <li>• <b>static</b>—Power allocated for interfaces is based on the class of powered device connected.</li> <li>• <b>high-power</b>—Power allocated for interfaces up to 65 W per port.</li> </ul>

### 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 4 on page 31](#) lists the classes and their power ratings as specified by the IEEE standards.

Table 4: 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 5 on page 32](#) shows the PoE configuration options and their default settings for the PoE controller and for the PoE interfaces.

Table 5: PoE Configuration Options and Default Settings

Option	Default	Description
PoE Controller Options		
<b><i>guard-band</i></b>	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.
<b><i>management</i></b>	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"> <li>• <b>class</b>—Power allocated for each interface can be configured.</li> <li>• <b>static</b>—Power allocated for interfaces is based on the class of powered device connected.</li> <li>• <b>high-power</b>—Power allocated for interfaces up to 65 W per port.</li> </ul>
Interface Options		
<b><i>disable (Power over Ethernet)</i></b>	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.
<b><i>priority (Power over Ethernet)</i></b>	low	Sets an interface's power priority to either <b>low</b> or <b>high</b> . 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.
<b><i>telemetries</i></b>	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 <b><i>interval (Power over Ethernet)</i></b> or <b><i>duration</i></b> .

## 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 | 36](#)

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

Table 6: 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

To configure PoE on an ACX2000 router:

#### CLI Quick Configuration

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

```
set poe management high-power guard-band 19
```

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

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *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 | 36](#)
- [Verifying the Telemetry Data \(History\) for the Specified Interface | 37](#)
- [Verifying PoE Global Parameters | 38](#)

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 index	Maximum power	Power consumption	Guard band	Management	Status	Lldp 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 | 30](#)

Example: Disabling a PoE Interface on ACX2000 Routers

IN THIS SECTION

- Requirements | 38
- Overview | 39
- Configuration | 39
- Verification | 39

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

Requirements

Before you begin:

- Configure PoE on all interfaces. See [“Example: Configuring PoE on ACX2000 Routers”](#) on page 33.

### 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

#### Step-by-Step Procedure

- Disable PoE on all interfaces.

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

- Disable PoE on a specific interface.

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

### 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 on EX Series Switches](#)

Troubleshooting PoE Interfaces on ACX2000 Universal Metro Routers

Problem

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

Solution

Check for the items shown in [Table 7 on page 40](#).

Table 7: Troubleshooting a PoE Interface

Items to Check	Explanation
Is interface PoE enabled?	Only interfaces <b>ge-0/1/3</b> and <b>ge-0/1/7</b> can function as PoE ports.
Has PoE capability been disabled for that interface?	Use the <b>show poe interface</b> 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 <b>show poe interface</b> command to check the maximum power provided by the interface.
If the <b>telemetries</b> option has been enabled for the interface, check the history of power consumption.	Use the <b>show poe telemetries</b> command to display the history of power consumption.

SEE ALSO

| [Monitoring and Troubleshooting PoE](#)  
| [Troubleshooting PoE Interfaces](#)

RELATED DOCUMENTATION

| [Configuring PoE on EX Series Switches](#)  
| [Monitoring and Troubleshooting PoE](#)  
| [Troubleshooting PoE Interfaces](#)

## Point-to-Point Protocol over Ethernet (PPPoE)

### IN THIS SECTION

- [PPPoE Overview | 41](#)
- [Configuring PPPoE | 45](#)
- [Disabling the Sending of PPPoE Keepalive Messages | 52](#)
- [Verifying a PPPoE Configuration | 53](#)
- [Tracing PPPoE Operations | 53](#)

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 | 42](#)
- [PPPoE Stages | 43](#)
- [Optional CHAP Authentication | 44](#)

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 43](#) 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

<i>Configuring the PPP Challenge Handshake Authentication Protocol</i>
<i>Evaluation Order for Matching Client Information in PPPoE Service Name Tables</i>
<i>Benefits of Configuring PPPoE Service Name Tables</i>
<a href="#">Configuring PPPoE   45</a>
<a href="#">Disabling the Sending of PPPoE Keepalive Messages   52</a>
<i>Configuring PPPoE Service Name Tables</i>
<i>Creating a Service Name Table</i>
<i>Configuring the Action Taken When the Client Request Includes an Empty Service Name Tag</i>
<i>Configuring the Action Taken for the Any Service</i>
<i>Assigning a Service to a Service Name Table and Configuring the Action Taken When the Client Request Includes a Non-zero Service Name Tag</i>
<i>Assigning an ACI/ARI Pair to a Service Name and Configuring the Action Taken When the Client Request Includes ACI/ARI Information</i>
<i>Limiting the Number of Active PPPoE Sessions Established with a Specified Service Name</i>
<i>Reserving a Static PPPoE Interface for Exclusive Use by a PPPoE Client</i>
<i>Enabling Advertisement of Named Services in PADO Control Packets</i>
<i>Assigning a Service Name Table to a PPPoE Underlying Interface</i>
<i>Example: Configuring a PPPoE Service Name Table</i>
<a href="#">Tracing PPPoE Operations   53</a>
<i>Troubleshooting PPPoE Service Name Tables</i>
<a href="#">Verifying a PPPoE Configuration   53</a>
<i>Ethernet Interfaces User Guide for Routing Devices</i>

## Configuring PPPoE

### IN THIS SECTION

- [Overview | 46](#)
- [Setting the Appropriate Encapsulation on the PPPoE Interface | 47](#)
- [Configuring PPPoE Encapsulation on an Ethernet Interface | 48](#)
- [Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface | 48](#)

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

### Overview

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

1. Configure PPPoE encapsulation for an Ethernet interface.
2. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
3. Optionally, configure the maximum transmission unit (MTU) of the interface.
4. Configure the operational mode as server.
5. Configure the PPPoE interface address.
6. Configure the destination PPPoE interface address.
7. Optionally, configure the MTU size for the protocol family.
8. Starting in Junos OS Release 10.0, optionally, configure one or more PPPoE service name tables and the action taken for each service in the tables.
9. Starting in Junos OS Release 12.3, optionally, disable the sending of PADS messages that contain certain error tags.

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

### ***Setting the Appropriate Encapsulation on the PPPoE Interface***

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

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

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

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

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

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

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

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

You can include this statement at the following hierarchy levels:

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

Perform the task appropriate for the interface on which you are using PPPoE. For more information on how to configure PPOE encapsulation on an ethernet interface and on an ATM-over-ADSL interface, see

[“Configuring PPPoE Encapsulation on an Ethernet Interface” on page 48](#) and [“Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface” on page 48](#).

### **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 | 3](#)

[Initial Configuration of Ethernet Interfaces | 4](#)

## 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 | 41](#)

*Ethernet Interfaces User Guide for Routing Devices*

## Verifying a PPPoE Configuration

### 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 | 41](#)

*Ethernet Interfaces User Guide for Routing Devices*

## Tracing PPPoE Operations

### IN THIS SECTION

- [Configuring the PPPoE Trace Log Filename | 54](#)
- [Configuring the Number and Size of PPPoE Log Files | 55](#)
- [Configuring Access to the PPPoE Log File | 55](#)
- [Configuring a Regular Expression for PPPoE Lines to Be Logged | 55](#)
- [Configuring the PPPoE Tracing Flags | 56](#)

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 | 53](#)[traceoptions \(PPPoE\) | 1400](#)

---

### **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 | 53](#)[traceoptions \(PPPoE\) | 1400](#)

---

### **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 | 53](#)[traceoptions \(PPPoE\) | 1400](#)

---

### **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 | 53](#)

[traceoptions \(PPPoE\) | 1400](#)

### ***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 | 53](#)

SEE ALSO

[PPPoE Overview | 41](#)

# Configuring Aggregated Ethernet Interfaces

## IN THIS CHAPTER

- [Configuring Aggregated Ethernet Interfaces and LACP | 57](#)
- [Link Protection of Aggregated Ethernet Interfaces | 133](#)
- [Scheduling on Aggregated Ethernet Interfaces | 138](#)
- [Load Balancing on Aggregated Ethernet Interfaces | 140](#)
- [Performance Monitoring on Aggregated Ethernet Interfaces | 179](#)
- [Periodic Packet Management | 183](#)
- [Understanding Ethernet Link Aggregation on ACX Series Routers | 188](#)

## Configuring Aggregated Ethernet Interfaces and LACP

## IN THIS SECTION

- [Aggregated Ethernet Interfaces Overview | 58](#)
- [Configuring an Aggregated Ethernet Interface | 64](#)
- [Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers | 65](#)
- [Configuring LACP for Aggregated Ethernet Interfaces | 67](#)
- [Configuring Junos OS for Supporting Aggregated Devices | 77](#)
- [Configuring the Number of Aggregated Ethernet Interfaces on the Device | 81](#)
- [Configuring Aggregated Ethernet Link Speed | 81](#)
- [Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces | 84](#)
- [Example: Configuring Aggregated Ethernet Interfaces | 84](#)
- [Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles | 86](#)
- [Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles | 94](#)
- [Configuring Tagged Aggregated Ethernet Interfaces | 98](#)
- [Configuring Untagged Aggregated Ethernet Interfaces | 99](#)
- [Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 100](#)

- [Understanding Independent Micro BFD Sessions for LAG | 113](#)
- [Example: Configuring Independent Micro BFD Sessions for LAG | 116](#)
- [Deleting an Aggregated Ethernet Interface | 127](#)

You can use Ethernet link aggregation to increase bandwidth and resiliency of Ethernet links by bundling or combining multiple full-duplex Ethernet links into a single virtual link. This virtual link interface is known as an aggregated Ethernet Interface. This topic provides an overview of the aggregated Ethernet Interface and explains how to configure an aggregated Ethernet Interface and its various features.

## Aggregated Ethernet Interfaces Overview

### IN THIS SECTION

- [Platform Support for Aggregated Ethernet Interfaces | 58](#)
- [Enhanced LAG Support on MX Series Routers | 60](#)
- [Enhanced LAG Support on PTX Series Routers | 61](#)
- [Configuration Guidelines for Aggregated Ethernet Interfaces | 61](#)

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. The Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-flow load balancing.

**NOTE:** For information about configuring circuit cross-connects over aggregated Ethernet, see *Circuit and Translational Cross-Connects Overview*.

For information about mixed rates and mixed modes on an aggregated Ethernet bundle, see [“Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles” on page 86](#).

### ***Platform Support for Aggregated Ethernet Interfaces***

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device and then associating a set of ports that have the same speed and are in full-duplex mode. The physical interfaces can be Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, Gigabit Ethernet IQ, 10-Gigabit



Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, or 10-Gigabit Ethernet IQ2 and IQ2-E. Generally, you cannot use a combination of these interfaces within the same aggregated link; however, you can combine Gigabit Ethernet and Gigabit Ethernet IQ interfaces in a single aggregated Ethernet bundle.

Starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

**NOTE:**

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

**TIP:**

Going forward:

- Aggregated Ethernet link with member links of different modes will be referred as *10-Gigabit Ethernet mixed mode aggregated Ethernet link*.
- Aggregated Ethernet link with member links of different rates will be referred as *mixed rate aggregated Ethernet link*.
- These aggregated Ethernet links will generically be referred as *mixed aggregated Ethernet links*.

Table 8 on page 60 lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

Table 8: Platform Support Matrix for Mixed Aggregated Ethernet Bundles

| 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>                                                                                                                             |
| 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-1XGE-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> |

**Enhanced LAG Support on MX Series Routers**

Starting in Junos OS Release 14.2, you can configure an enhanced link aggregation group (LAG) on MX Series routers. 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.

In the absence of enhanced LAG support, 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 installation of 17 next hops per VLAN created. Thus the number of next hops supported on the routers with aggregated Ethernet interfaces is significantly reduced.

With the enhanced LAG support, when the `[edit chassis network-services enhanced-ip]` statement is configured, child next hops are not created for member links and, as a result, a higher number of next hops can be supported.

Note that the enhanced LAG feature is only supported when the router's network services is set to operate in the **enhanced-ip** mode. This feature is not supported when the router's network services is set to operate in the **enhanced-ethernet** mode.

### ***Enhanced LAG Support on PTX Series Routers***

Starting in Junos OS Release 18.1, Junos OS supports removal of child next hop usage for aggregated Ethernet Interfaces and clients on PTX Series routers with FPC3-PTX-U2 and FPC3-PTX-U3. Removal of child next hop usage helps reduce the memory and CPU resources required to support aggregated Ethernet Interfaces and improves the overall system performance and scaling numbers. This feature is enabled by default if the network services mode on the router is configured to **enhanced-mode**. You can disable this feature by using the `set chassis aggregated-devices disable-lag-enhanced`. You must reboot the router for the configuration to take effect.

Previously, each unicast next hop over aggregated Ethernet Interfaces resulted in creation of a number of children next hops as well. For an aggregated Ethernet Interface with 16 member links, addition of one unicast next hop over the aggregated Ethernet Interface results in installing total of 17 next hops. As a result, with aggregated Ethernet configuration, the number of next hops supported decreases in proportion to the number of aggregated Ethernet links.

**NOTE:** Child next hop optimizations are supported for aggregated Ethernet Interfaces, Interfaces that make use of aggregated Ethernet Interfaces, and for both unicast and multicast scenarios. For more information, see [“Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers” on page 65](#).

### ***Configuration Guidelines for Aggregated Ethernet Interfaces***

- For Junos OS Evolved, when a new interface is added as a member to the aggregated Ethernet bundle, a link flap event is generated. When you add an interface to the bundle, 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.
- For Junos OS Evolved, there is no restriction on the maximum number of AE interfaces in a mixed-rate AE bundle. Because all child logical interfaces belong to same AE physical interface and share the same selector, using much less load balance memory, mixed-rate AE interface configurations should go through even if they exceed 64 logical interfaces.
- Aggregated Ethernet for subscriber management supports only **gigether-options** on the member link interfaces. If you configure **ether-options** on the member link interfaces, subscriber management does not work properly. The **ether-options** configuration will cause issues with subscriber accounting and statistics.

- Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, PICs, or MPCs.
- All Juniper routers support at least eight physical interfaces per aggregated Ethernet bundle. See [maximum-links](#) configuration page for platform specific limits.
- On M Series and T Series routers, you can create a maximum of 1024 logical interfaces on an aggregated Ethernet interface.
- Simple filters are not supported for interfaces in aggregated Ethernet bundles:
  - On M Series routers, simple filters are supported in Gigabit Ethernet Enhanced Intelligent Queuing interfaces only, except when the interface is part of an aggregated Ethernet bundle.
  - On MX Series routers, simple filters are supported in Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only, except when the interface is part of an aggregated Ethernet bundle.

For more information about simple filters, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

- On the aggregated Ethernet bundle, no IQ-specific capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available. For more information about IQ-specific capabilities, see [“Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs”](#) on page 291.
- Aggregated Ethernet interfaces can be either tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on MX Series routers support the configuration of **flexible-vlan-tagging** and **native-vlan-id** on dual-tagged frames, which consist of the following configuration statements:
  - **inner-tag-protocol-id**
  - **inner-vlan-id**
  - **pop-pop**
  - **pop-swap**
  - **push-push**
  - **swap-push**
  - **swap-swap**

In all cases, you must set the number of aggregated Ethernet interfaces on the chassis. You can also set the link speed and the minimum links in a bundle.

For more information on tagging in PTX Series routers, see [“Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers”](#) on page 65.

- When configuring mixed aggregated Ethernet bundles on T640, T1600, T4000, and TX Matrix Plus routers, consider the following:
  - A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
  - Link Aggregation Control Protocol (LACP), aggregated Ethernet link protection, and LACP link protection are supported only on mixed aggregated Ethernet link configured on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).

- 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. [Table 8 on page 60](#) lists the platforms and corresponding hardware components 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 configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4), ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.
- When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC with CFP, changing aggregated Ethernet link protection or LACP 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 with CFP, 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.
- Use the **show interfaces aggregate-interface extensive** and **show interfaces aggregate.logical-interface** commands to show the bandwidth of the aggregate. Also, the SNMP object identifier **ifSpeed/ifHighSpeed** shows the corresponding bandwidth on the aggregate logical interface if it is configured properly.

SEE ALSO

| [Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs](#) | 291

## Configuring an Aggregated Ethernet Interface

On Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on M Series and T Series routers, you can associate a physical interface with an aggregated Ethernet interface.

**NOTE:** On a Junos Fusion, you can include extended ports (physical interface on a satellite device that provides a connection to servers or endpoints) or local ports in link aggregation groups (LAGs) and MC-LAGs, but not both. For information on extended ports, see *Understanding Junos Fusion Ports*.

To configure an aggregated Ethernet interface:

1. Specify that you want to configure the link aggregation group interface.

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

2. Configure the aggregated Ethernet interface.

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

You specify the interface instance number *x* to complete the link association; *x* can be from 0 through 127, for a total of 128 aggregated interfaces on M Series and T Series routers and can be from 1 through 480, for a total of 480 aggregated interfaces on MX Series routers. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. Aggregated interfaces are numbered from **ae0** through **ae4092**.

**NOTE:** On MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.

You must also include a statement defining **aex** at the **[edit interfaces]** hierarchy level. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see [“Ethernet Interfaces Overview” on page 3](#), and for a sample configuration, see [“Example: Configuring Aggregated Ethernet Interfaces” on page 84](#).

**NOTE:** In general, aggregated Ethernet bundles support the features available on all supported interfaces that can become a member link within the bundle. As an exception, Gigabit Ethernet IQ features and some newer Gigabit Ethernet features are not supported in aggregated Ethernet bundles.

Gigabit Ethernet IQ and SFP interfaces can be member links, but IQ- and SFP-specific features are not supported on the aggregated Ethernet bundle even if all the member links individually support those features.

You need to configure the correct link speed for the aggregated Ethernet interface to eliminate any warning message.

**NOTE:** Before you commit an aggregated Ethernet configuration, ensure that link mode is not configured on any member interface of the aggregated Ethernet bundle; otherwise, the configuration commit check fails.

## Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers

### IN THIS SECTION

- [Support for Aggregated Ethernet Interfaces | 66](#)
- [Understanding Aggregated Ethernet Interfaces on Junos OS Evolved | 66](#)
- [How to Configure Aggregated Ethernet Interfaces | 66](#)

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a link aggregation group (LAG) or bundle. Link aggregation can be used for point-to-point connections. It 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.

This topic describes how to configure aggregated Ethernet interfaces on PTX Series Packet Transport Routers.

### ***Support for Aggregated Ethernet Interfaces***

On PTX Series Packet Transport Routers, aggregated Ethernet support includes the following features:

- A consistent interface type (**et fpc/pic/port**) across all Ethernet interfaces.
- Ability to bundle multiple Ethernet interfaces
- Ability to bundle mixed-rate links on the same aggregated Ethernet interface
- Fault tolerance
- Load balancing between child links
- Advanced features including flexible VLAN tagging and Ethernet services encapsulation

### ***Understanding Aggregated Ethernet Interfaces on Junos OS Evolved***

In PTX Series routers running Junos OS Evolved such as PTX1003 Series routers, VLAN-tagged aggregated Ethernet interfaces behave differently than interfaces that are not tagged.

- In an untagged LAG, child IFL members are created. Requests are made per child IFL member. The results are aggregated and displayed in the CLI.
- In a VLAN-tagged LAG, extra child IFLs are not created as part of the aggregated Ethernet bundle. Link IFL statistics and marker statistics for child IFLs are not displayed.

### ***How to Configure Aggregated Ethernet Interfaces***

Aggregated Ethernet interfaces can use interfaces from different FPCs or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
```

```
[edit interfaces]
  et-0/0/0 {
    together-options {
      802.3ad ae0;
    }
  }
  et-0/0/1 {
    together-options {
      802.3ad ae0;
    }
  }
```



```

}
ae0 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 200.200.1.2/24;
    }
  }
  unit 1 {
    vlan-id 101;
    family inet {
      address 200.200.2.2/24;
    }
  }
}

```

## Configuring LACP for Aggregated Ethernet Interfaces

### IN THIS SECTION

- [Configuring the LACP Interval | 70](#)
- [Configuring LACP Link Protection | 70](#)
- [Configuring LACP Hold-Up Timer to Prevent Link Flapping on LAG Interfaces | 74](#)
- [Tracing LACP Operations | 75](#)
- [Sample Configuration for Configuring Aggregated Ethernet LACP on Tagged and Untagged Interfaces | 75](#)

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.

**NOTE:** Starting with Junos OS Release 14.1, you can configure aggregated Ethernet interfaces with LACP on logical systems within an MX Series router.

For Multichassis Link Aggregation (MC-LAG), you must specify the **system-id** and **admin key**. MC-LAG peers use the same **system-id** while sending the LACP messages. The **system-id** can be configured on the MC-LAG network device and synchronized between peers for validation.

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 was 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.

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.

**NOTE:** LACP can link together multiple different physical interfaces, but only features that are supported across all of the linked devices will be 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 will only support up to 8 forwarding classes. Similarly, linking together a PIC that supports WRED with a PIC that does not support it will result in a LAG bundle that does not support WRED.

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.

If you restart the Link Aggregation Control Protocol (LACP) process consecutively without adequate sleep or pause time between the successive restarts, the LACP links might flap. To avoid this problem, you must restart the LACP process only after the refresh time of the periodic packet management (PPM) process is completed.

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;
}
```

Starting with Junos OS Release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.

To override the IEEE 802.3ad standard and facilitate subsecond failover, include the **fast-failover** statement at the **[edit interfaces *interface-name* aggregated-ether-options lACP]** hierarchy level.

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 use 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 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.

For more information, see the following sections:

### Configuring the LACP Interval

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

```
[edit interfaces interface-name aggregated-ether-options lacp]
periodic interval;
```

The interval can be fast (every second) or slow (every 30 seconds). You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.

**NOTE:** Starting with Junos OS Release 11.4, source address filtering does not work when LACP is enabled. This behavior is not applicable to T Series routers and PTX Series Packet Transport Routers. For more information about source address filtering, see [“Configuring MAC Address Filtering for Ethernet Interfaces” on page 21](#).

Percentage policers are not supported on aggregated Ethernet interfaces with the CCC protocol family configured. For more information about percentage policers, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

Generally, LACP is supported on all untagged aggregated Ethernet interfaces. For more information, see [“Configuring Untagged Aggregated Ethernet Interfaces” on page 99](#).

For M Series Multiservice Edge Routers with enhanced Flexible PIC Concentrators (FPCs) and T Series routers, LACP over VLAN-tagged aggregated Ethernet interfaces is supported. For 8-port, 12-port, and 48-port Fast Ethernet PICs, LACP over VLAN-tagged interfaces is not supported.

LACP **Fast Periodic**, which is achieved by configuring fast (every second) intervals for periodic transmission of LACP packets, is supported with graceful Routing Engine switchover (GRES) on MX Series routers only.

### Configuring LACP Link Protection

#### IN THIS SECTION

- [Enabling LACP Link Protection | 71](#)
- [Configuring LACP System Priority | 72](#)
- [Configuring LACP System Identifier | 72](#)

- [Configuring LACP administrative Key | 73](#)
- [Configuring LACP Port Priority | 73](#)

To force active and standby links within an aggregated Ethernet, you can configure LACP link protection and system priority at the aggregated Ethernet interface level using the **link-protection** and **system-priority** statements. Configuring values at this level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[gigether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

**NOTE:**

- When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.
- LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.

### ***Enabling LACP Link Protection***

To enable LACP link protection for an aggregated Ethernet interface, use the **link-protection** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
link-protection;
  disable;
  revertive;
  non-revertive;
}
```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, once a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch and the current link remains active.

If LACP link protection is configured to be nonrevertive at the global (**[edit chassis]** hierarchy) level, you can add the **revertive** statement to the LACP link protection configuration to override the nonrevertive setting for the interface. In revertive mode, the addition of a higher-priority link to the aggregator results in LACP performing a priority recalculation and switching from the current active link to the new active link.



**CAUTION:** If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

We strongly recommend that you use LACP on both ends of the aggregator, when you connect an aggregated Ethernet interface with two member interfaces of MX Series routers to any other vendor device. Otherwise, the vendor device (say a Layer 2 switch, or a router) will not be able to manage the traffic coming from the two link aggregated Ethernet bundle. As a result, you might observe the vendor device sending back the traffic to the backup member link of the aggregated Ethernet interface.

Currently, MX-MPC2-3D, MX-MPC2-3D-Q, MX-MPC2-3D-EQ, MX-MPC1-3D, MX-MPC1-3D-Q, and MPC-3D-16XGE-SFP do not drop traffic coming back to the backup link, whereas DPCE-R-Q-20GE-2XGE, DPCE-R-Q-20GE-SFP, DPCE-R-Q-40GE-SFP, DPCE-R-Q-4XGE-XFP, DPCE-X-Q-40GE-SFP, and DPCE-X-Q-4XGE-XFP drop traffic coming to the backup link.

### Configuring LACP System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the interface, use the **system-priority** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 to 65,535.

### Configuring LACP System Identifier

To configure the LACP system identifier for aggregated Ethernet interfaces, use the **system-id** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
system-id system-id;
```

Starting with Junos OS Release 13.3, you must not configure the LACP system identifier by using the **system-id system-id** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level to be all zeros (00:00:00:00:00:00). If you attempt to commit a configuration with the system identifier to be all zeros, an error occurs during the commit operation.

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.

### Configuring LACP administrative Key

To configure an administrative key for LACP, include the **admin-key number** statement at the **[edit interfaces aex aggregated-ether-options lacp]** hierarchy level:

```
[edit interfaces ae x aggregated-ether-options-lacp]
admin-key number;
```

**NOTE:** You must configure MC-LAG to configure the **admin-key** statement. For more information about MC-LAG, see *Configuring Multichassis Link Aggregation on MX Series Routers*.

### Configuring LACP Port Priority

To configure LACP port priority for aggregated Ethernet interfaces, use the **port-priority** statement at the **[edit interfaces interface-name together-options 802.3ad aeX lacp]** or **[edit interfaces interface-name fastether-options 802.3ad aeX lacp]** hierarchy levels:

```
[edit interfaces interface-name together-options 802.3ad aeX lacp]
port-priority priority;
```

The port priority is a 2-octet field that is part of the LACP port ID. The LACP port ID consists of the port priority as the two most-significant octets and the port number as the two least-significant octets. The system with the numerically lower value for port priority has the higher priority. By default, port priority is 127, with a range of 0 to 65,535.

Port aggregation selection is made by each system based on the highest port priority and is assigned by the system with the highest priority. Ports are selected and assigned starting with the highest priority port of the highest priority system and working down in priority from there.

**NOTE:** Starting with Junos OS Release 9.3, port aggregation selection (discussed previously) is performed for the active link when LACP link protection is enabled. Without LACP link protection, port priority is not used in port aggregation selection.

### ***Configuring LACP Hold-Up Timer to Prevent Link Flapping on LAG Interfaces***

On link aggregation group (LAG) interfaces, when a member (child) link goes down, its state changes from current to expired. This link might flap from the current state to the expired state and back to current state when it receives intermittent LACP protocol data units (PDUs) and keepalive timeouts. Such flapping can adversely affect the traffic on the link.

To prevent excessive flapping of a LAG child link, you can configure a hold-up timer on the LAG interface that is applicable to all member links on that particular interface. To hold up, in networking terms, means to prevent the transitioning of an interface from down to up for a specified time interval.

When configured, the hold-up timer is triggered when an LACP state machine tries to move to the current state from the expired or default state when it receives an LACP PDU. The hold-up timer is triggered only if the LACP state machine had acquired the current state at least once earlier. The timer is not triggered if LACP attempts to transition to the current state for the first time. LACP monitors the PDUs received on the child link but prevents the link from transitioning to current state. If no flapping is observed when the link receives the PDUs, the hold-up timer expires and triggers the member link to transition back to the current state. This transition is triggered as soon as the hold-up timer expires and not necessarily when the link receives a PDU.

To configure LACP hold-up timer for LAG interface, use the **hold-time up** statement at the **[edit interfaces aex aggregated-ether-options lacp]** hierarchy level.



**NOTE:**

- The hold-up timer keeps running even when the interface that receives the LACP PDU moves to the port disable state. The timer is then restarted if, before the timer expires, the interface comes up again and receives an LACP PDU from its neighbor. This ensures that the timer is maintained even during a quick physical port flap.
- When the following events occur, a hold-up timer is not triggered until the member link acquires the current state after the event:
  - LACP daemon restart
  - Deactivation and reactivation of child or aggregated Ethernet interface
  - Deletion and reconfiguration of child or aggregated Ethernet interface
  - System reboot
  - Routing Engine switchover

**Tracing LACP Operations**

To trace the operations of the LACP process, include the **traceoptions** statement at the **[edit protocols lacp]** hierarchy level:

```
[edit protocols lacp]
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

**Sample Configuration for Configuring Aggregated Ethernet LACP on Tagged and Untagged Interfaces**

Following configurations are examples of configuring aggregated Ethernet LACP on VLAN-tagged and untagged interfaces:

**LACP with VLAN-Tagged Aggregated Ethernet**

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
```

```

ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}

```

### LACP with Untagged Aggregated Ethernet

```

[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 10.1.1.4;
          priority 200;
        }
      }
    }
  }
}

```

```

    }
  }
}
}
}

```

## SEE ALSO

*Junos OS Administration Library*

## Configuring Junos OS for Supporting Aggregated Devices

Junos OS supports the aggregation of physical devices into defined virtual links, such as the link aggregation of Ethernet interfaces defined by the IEEE 802.3ad standard.

Tasks for configuring aggregated devices are:

- [Configuring Virtual Links for Aggregated Devices | 77](#)
- [Configuring LACP Link Protection at the Chassis Level | 78](#)
- [Enabling LACP Link Protection | 79](#)
- [Configuring System Priority | 80](#)
- [Configuring the Maximum Links Limit | 80](#)
- [Configuring PPM on Junos Fusion | 80](#)

### Configuring Virtual Links for Aggregated Devices

To define virtual links, you need to specify the associations between physical and logical devices within the **[edit interfaces]** hierarchy, and assign the correct number of logical devices by including the **device-count** statement at the **[edit chassis aggregated-devices ethernet]** and **[edit chassis aggregated-devices sonet]** hierarchy levels:

```

[edit chassis]
aggregated-devices {
  ethernet {
    device-count number;
  }
  sonet {
    device-count number;
  }
}

```

```
}
}
```

The aggregated interfaces are numbered from **ae0** through **ae4091**. The maximum number of aggregated interfaces supported by different routers is listed below:

- For PTX Series routers, you can configure a maximum of 128 aggregated interfaces.
- For M Series and T Series routers, you can configure a maximum of 128 aggregated interfaces (LAG bundles).
- In Junos release 14.2R2 and earlier, you can configure a maximum of 480 aggregated interfaces on MX Series routers.
- In Junos release 14.2R3 and later, you can configure a maximum of 1000 aggregated interfaces on MX240, MX480, and MX960 routers.
- In Junos release 14.2R3 and later, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.
- In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 480 aggregated interfaces on MX240, MX480, and MX960 routers.
- In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.

For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 64, numbered from **as0** through **as63**. In releases before Junos OS Release 13.2, the maximum was 16.

### **Configuring LACP Link Protection at the Chassis Level**

Link Aggregation Control Protocol (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. 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 link protection enables you to force active and standby links within an aggregated Ethernet. You configure LACP link protection by using the **link-protection** and **system-priority** statements at either the chassis or interface level and by configuring port priority at the interface level using the **system-priority** statement. Configuring LACP parameters at the chassis level results in all aggregated Ethernet interfaces using the defined values unless overridden by the LACP configuration on a specific interface.

```
[edit chassis]
aggregated-devices {
  ethernet {
    lacp {
      link-protection {
```

```

        non-revertive;
    }
    system-priority priority;
}
}
}

```

**NOTE:** LACP link protection also uses port priority. You can configure port priority at the Ethernet interface **[gigether-options]** hierarchy level using the **port-priority** statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).

#### SEE ALSO

##### *Enabling LACP Link Protection*

To enable LACP link protection for aggregated Ethernet interfaces on the chassis, use the **link-protection** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```

[edit chassis aggregated-devices ethernet lacp]
link-protection {
    non-revertive;
}

```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, after a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch, and the current link remains active.

**BEST PRACTICE:** (MX Series) 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.



**CAUTION:** If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

### Configuring System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the chassis, use the **system-priority** statement at the **[edit chassis aggregated-devices ethernet lacp]** hierarchy level:

```
[edit chassis aggregated-devices ethernet lacp]
system-priority priority;
```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 through 65,535.

### Configuring the Maximum Links Limit

To configure the maximum links limit, use the **maximum-links** statement at the **[edit chassis aggregated-devices]** hierarchy level:

```
[edit chassis aggregated-devices]
maximum-links maximum-links-limit;
```

### Configuring PPM on Junos Fusion

If you use Junos Fusion with Junos OS Release 14.2R3, you need to ensure that link aggregation (and STP) work properly by configuring timers for the periodic packet management (PPM) daemons on the aggregation and satellite devices. We recommend using the following timer values:

```
[edit routing-options ppm]
redistribution-timer 120;
tcp-keepalive-interval 3000;
tcp-keepalive-idle 3000;
```

Starting in Junos OS Release 14.2R4, the timer values that ensure proper link aggregation and STP functions are configured by default if you use Junos Fusion with Junos OS.

SEE ALSO

| *Configuring Aggregated SONET/SDH Interfaces*

## Configuring the Number of Aggregated Ethernet Interfaces on the Device

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them.

For M Series and T Series routers you can configure a maximum of 128 aggregated interfaces (LAG bundles). On MX Series routers running Junos release 14.2R2 and earlier, you can configure a maximum of 480 aggregated interfaces. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. For MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces. In all cases the aggregated interfaces are numbered from **ae0** through **ae4092**.

**NOTE:** On a Junos Fusion Fabric, you can include extended ports (physical interface on a satellite device that provides a connection to servers or endpoints) or local ports in link aggregation groups (LAGs) and MC-LAGs, but not both. For information on extended ports, see *Understanding Junos Fusion Ports*.

For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 16, numbered from **as0** through **as15**.

1. Specify that you want to access the aggregated Ethernet configuration on the device.

```
user@host# edit chassis aggregated-devices ethernet
```

2. Set the number of aggregated Ethernet interfaces.

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

You must also specify the constituent physical links by including the **802.3ad** statement at the **[edit interfaces interface-name fastether-options]** or **[edit interfaces interface-name gigether-options]** hierarchy level.

## Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, there are exceptions.

Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers. For example, these mixes are supported:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.

Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.

#### NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of mixed rates and mixed modes on T640, T1600, T4000, TX Matrix Plus, and PTX routers, you need to configure the **mixed** option for the **[edit interfaces *aex* aggregated-ether-options link-speed]** statement.

To set the required link speed:

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

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

2. Configure the link speed.

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

***speed*** can be 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 interfaces 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.
- **50g**—Links are 50 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**—Links are of various speeds.
- **oc192**—Links are OC192.

SEE ALSO

| [aggregated-ether-options](#) | **1108**

## Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces

T Series and TX Matrix routers support multicast statistics collection on aggregated 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 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:

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

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

```
[edit interfaces]
ae0 {
    multicast-statistics;
}
```

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

SEE ALSO

[multicast-statistics](#) | 1280

[Configuring Multicast Statistics Collection on Ethernet Interfaces](#) | 10

### Example: Configuring Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit chassis]
aggregated-devices {
    ethernet {
        device-count 15;
```

```
}  
}
```

```
[edit interfaces]  
ge-1/3/0 {  
  gicether-options {  
    802.3ad ae0;  
  }  
}  
ge-2/0/1 {  
  gicether-options {  
    802.3ad ae0;  
  }  
}  
ae0 {  
  aggregated-ether-options {  
    link-speed 1g;  
    minimum-links 1;  
  }  
}  
vlan-tagging;  
unit 0 {  
  vlan-id 1;  
  family inet {  
    address 10.0.0.1/24;  
  }  
}  
unit 1 {  
  vlan-id 1024;  
  family inet {  
    address 10.0.0.2/24;  
  }  
}  
unit 2 {  
  vlan-id 1025;  
  family inet {  
    address 10.0.0.3/24;  
  }  
}  
unit 3 {  
  vlan-id 4094;  
  family inet {  
    address 10.0.0.4/24;  
  }  
}
```

```
}
}
```

## SEE ALSO

[Configure 'link-speed' for Gigabit Ethernet based Aggregated Ethernet interface bundles](#)

## Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles

### IN THIS SECTION

- [Aggregated Ethernet Bundle with Mixed Rates and Mixed Modes on T Series Routers | 86](#)
- [Aggregated Ethernet Bundles with Mixed Rates on MX Series Routers and PTX Series Routers | 90](#)

You can configure the member links of an aggregated Ethernet bundle with any combination of rates—also known as mixed rates—on T Series, MX Series, and PTX Series routers. The bandwidth that is provided by an aggregated Ethernet bundle can be utilized completely and efficiently when the links are configured with different rates.

You can configure mixed modes on T Series routers. In a mixed-mode configuration, the member links of an aggregated Ethernet bundle are configured in LAN mode as well as in WAN mode for 10-Gigabit Ethernet interfaces. For information on the first Junos OS release that supports aggregated Ethernet bundles on the various Juniper Networks routers, see [Feature Explorer](#).

The following sections provide an overview about mixed rates and mixed modes on various platforms:

### ***Aggregated Ethernet Bundle with Mixed Rates and Mixed Modes on T Series Routers***

#### IN THIS SECTION

- [Understanding Mixed Rates and Mixed Modes | 87](#)
- [Platform Support Matrix for Mixed Aggregated Ethernet Bundles | 87](#)
- [Guidelines to Follow When Configuring Aggregated Ethernet Bundles with Mixed Rates and Mixed Modes | 88](#)

The following sections explain mixed rates and mixed modes on T Series routers:

### ***Understanding Mixed Rates and Mixed Modes***

Starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

#### **NOTE:**

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well.

In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

### **Going forward:**

- An aggregated Ethernet link with member links of different modes is referred to as *10-Gigabit Ethernet mixed-mode aggregated Ethernet link*.
- An aggregated Ethernet link with member links of different rates is referred to as *mixed-rate aggregated Ethernet link*.
- These aggregated Ethernet links will generically be referred to as *mixed aggregated Ethernet links*.

### ***Platform Support Matrix for Mixed Aggregated Ethernet Bundles***

[Table 9 on page 88](#) lists the platforms and corresponding hardware components that support mixed aggregated Ethernet bundles.

**Table 9: Platform Support Matrix for Mixed Aggregated Ethernet Bundles**

| 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>                                                                                                                             |
| 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> |

**Guidelines to Follow When Configuring Aggregated Ethernet Bundles with Mixed Rates and Mixed Modes**

In addition to the configuration guidelines for aggregated Ethernet interfaces in [“Aggregated Ethernet Interfaces Overview” on page 58](#), you must consider the following as well when configuring mixed modes and mixed rates on aggregated Ethernet bundles on T640, T1600, T4000, and TX Matrix Plus routers:

- A maximum of 16 member links can be configured to form a mixed aggregated Ethernet link.
- Link Aggregation Control Protocol (LACP), aggregated Ethernet link protection, and LACP link protection are supported only on mixed aggregated Ethernet bundles configured on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).

- 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. [Table 9 on page 88](#) lists the platforms and corresponding hardware components 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 configuring a mixed aggregated Ethernet link on a 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4), ensure that you add both the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP to the aggregated Ethernet bundle. Moreover, both these 50-Gigabit Ethernet interfaces must be included in the same aggregated Ethernet bundle.
- When a mixed aggregated Ethernet link is configured on a 100-Gigabit Ethernet PIC with CFP, changing aggregated Ethernet link protection or LACP 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 with CFP, 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.

## Aggregated Ethernet Bundles with Mixed Rates on MX Series Routers and PTX Series Routers

### IN THIS SECTION

- [Understanding Mixed Rates | 90](#)
- [Platform Support Matrix for Mixed Aggregated Ethernet Bundles on MX Series Routers | 91](#)
- [Supported Features | 92](#)

The following sections explain mixed rates on aggregated Ethernet bundles:

### **Understanding Mixed Rates**

Starting with Junos OS Release 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers, thereby enabling you to configure the member links with any combination of rates— 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet— on an aggregated Ethernet bundle.

**NOTE:** When you mix a 10-Gigabit Ethernet Interface in LAN mode and a 10-Gigabit Ethernet Interface in WAN mode in the same aggregate on MX-Series, it is not considered a mixed-rate aggregate. To mix the interfaces having 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.

You must configure the **mixed** statement explicitly at the **[edit interfaces *interface-name* aggregated-ether-options link-speed]** hierarchy level to:

- Enable the mixed-rate configuration on all the aggregated Ethernet bundles on the egress side of the Modular Port Concentrators (MPCs).
- Enable the router to detect child links of different speed values in the aggregated Ethernet bundle.

You can also configure the minimum bandwidth on an aggregated Ethernet bundle when you configure mixed rate on that aggregated Ethernet bundle.



**NOTE:** The **minimum-link** statement cannot be configured with mixed rates.

Mixed rates cannot be configured for aggregated Ethernet bundles on the egress side of the Dense Port concentrators (DPCs).

When you configure mixed rate on a homogeneous aggregated Ethernet bundle— where all the links in the bundle are of the same speed— the aggregated Ethernet bundle goes down and then comes up with the mixed-rate configuration.

### **Platform Support Matrix for Mixed Aggregated Ethernet Bundles on MX Series Routers**

Table 10 on page 91 lists the platforms and corresponding MPCs that support mixed aggregated Ethernet bundles on MX Series routers.

**Table 10: Platform Support Matrix for Mixed 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          |

Table 10: Platform Support Matrix for Mixed Aggregated Ethernet Bundles on MX Series Routers (*continued*)

| Supported MPCs            | Supported Platform                      | Initial Release |
|---------------------------|-----------------------------------------|-----------------|
| MPC7E 10G (MPC7E-10G)     | MX240, MX480, MX960, MX2010, and MX2020 | 15.1F5          |
| MPC8E (MX2K-MPC8E)        | MX2010 and MX2020                       | 15.1F5          |
| MPC9E (MX2K-MPC9E)        | MX2010 and MX2020                       | 15.1F5          |
| MPC10E (MPC10E-15C-MRATE) | MX240, MX480, and MX960                 | 19.1R1          |

**Supported Features**

The following features are supported on mixed-rate aggregated Ethernet bundles on MX Series routers and PTX Series Routers:

**NOTE:** Starting with Junos OS Release 19.2R1, you can configure any number of aggregated Ethernet Interfaces in a mixed rate aggregated Ethernet bundle on MX Series routers and PTX Series routers. In earlier releases, you can configure only 64 aggregated Ethernet Interfaces in a mixed-rate aggregated Ethernet bundle.

- Load balancing of traffic in proportion to the member-link speed.
- Non link-protect mode. For more information, see [“Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection” on page 139](#).
- LACP for slow and fast interval for periodic transmission of LACP packets.
- Port-based network access control (NAC).
- Scheduler parameters for aggregated interface member links in a scaled manner with the **member-link-scheduler scale** statement at the **[edit class-of-service interfaces]** hierarchy level.
- Layer 3 features only.
- Configuration of following statements as percentages only for mixed rates at the **[edit class-of-service schedulers scheduler-name]** hierarchy level:
  - **buffer-size**
  - **excess-rate**
  - **shaping-rate**
  - **transmit-rate**

- Configuration of the following statements for mixed rates at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level:
  - **drop-profile-map**
  - **excess-priority**
  - **priority**
  - **transmit-rate (rate-limit | exact)**
- The **shared-bandwidth-policer** statement at the **[edit firewall policer *policer-name*]** hierarchy level.
- The **scheduler-maps *map-name*** statement at the **[edit class-of-service]** hierarchy level.
- Unicast load balancing, where the load balancing happens on ingress-only selectors.
- Multicast load balancing
- Make-before-break (MBB) for multicast LDP (MLDP) and fast reroute (FRR).
- Source class usage (SCU) and destination class usage (DCU) accounting.
- Families **inet**, **inet6**, **mpls** and **iso**.
- Enhanced IP network services.
- LDP tunneling and OAM link fault management (LFM).

The following features are not supported on mixed-rate aggregated Ethernet bundles on MX Series routers and PTX Series routers:

- Adaptive load balancing
- Hierarchical schedulers on aggregated Ethernet bundles and the scheduling on logical interfaces (per-unit scheduling).
- Shaping rate, where traffic shaping is achieved by specifying the amount of bandwidth to be allocated to a logical interface.
- The **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level.
- Ingress queuing.
- Options that are configured with nonpercentage values at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level.
- The **member-link-scheduler replicate** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level.
- Aggregated Ethernet link protection and link protection on a 1:1 model.
- LACP link protection.

- Layer 2 features.
- The **target-routing-instance** (*routing-instance-name* | **default**) statement at the [edit access domain map *domain-map-name*] hierarchy level.

#### SEE ALSO

[minimum-bandwidth](#) | 1272

[P2-10G-40G-QSFPP PIC Overview](#) | 207

[Understanding the P2-100GE-OTN PIC](#) | 433

## Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles

### IN THIS SECTION

- [Configuring Mixed Rates and Mixed Modes on an Aggregated Ethernet Bundle on T Series Routers](#) | 95
- [Configuring Mixed Rates on Aggregated Ethernet Bundles on MX Series Routers](#) | 97

The following sections explain how to configure mixed rates and mixed modes on various platforms:

### **Configuring Mixed Rates and Mixed Modes on an Aggregated Ethernet Bundle on T Series Routers**

In releases before Junos OS Release 13.2, all interfaces that form an aggregated Ethernet bundle must have the same speed and must be in full-duplex mode. Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers. Following mixed rates and mixed modes are supported:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

#### **NOTE:**

- Member links of 50-Gigabit Ethernet can be configured using only the 50-Gigabit Ethernet interfaces of the 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of an aggregated Ethernet bundles in mixed rate or mixed mode on T640, T1600, T4000, and TX Matrix Plus routers:

1. Go to **[edit chassis]** hierarchy level.

```
user@host# edit chassis
```

2. Configure the number of aggregated logical devices available to the router:

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

3. Go to the **[edit interfaces]** hierarchy level.

```
user@host# top
user@host# edit interfaces
```

4. Configure the minimum number of links that are required for the aggregated Ethernet bundle to be labeled *up*:

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

**NOTE:** By default, only one link needs to be up for the bundle to be labeled *up*.

5. Configure the **link-speed** statement and specify the **mixed** option for the **link-speed** statement to indicate the mixed-rate and mixed-mode support for the aggregated Ethernet bundle configuration.

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

**NOTE:** It is mandatory to configure the **mixed** option for aggregated Ethernet bundles for the PD-1CE-CFP-FPC4 PIC.

On aggregated Ethernet bundles in MX Series routers, when the **mixed** statement at the **[edit interfaces aex aggregated-ether-options link-speed]** hierarchy level is not configured, the mixed rate configuration is applied by default.

6. Configure the members links of the aggregated Ethernet bundle.

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

7. Configure an interface family and an IP address for the aggregated Ethernet bundle.

```
[edit interfaces]
user@host# set aex unit number family (inet | inet6 | mpls ) address address
```

8. Commit the configuration.

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

### Configuring Mixed Rates on Aggregated Ethernet Bundles on MX Series Routers

Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers. You can now configure the member links with any combination of rates—10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet—on an aggregated Ethernet bundle.

To configure mixed rates on an aggregated Ethernet bundle on MX Series routers:

1. Go to the **[edit chassis]** hierarchy level.

```
user@host# edit chassis
```

2. Configure the number of aggregated logical devices available to the router.

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

3. Go to the **[edit interfaces]** hierarchy level.

```
user@host# top
user@host# edit interfaces
```

4. Configure the **link-speed** statement and specify the **mixed** option for the **link-speed** statement to indicate the mixed-rate support for the aggregated Ethernet bundle configuration.

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

5. Configure the members links of the aggregated Ethernet bundle:

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

6. Configure an interface family for the aggregated Ethernet bundle as inet, inet6, or mpls:

```
[edit interfaces]
user@host# set aex unit number family (inet | inet6 | mpls)
```

7. Configure the minimum bandwidth unit as **bps**, **gbps**, **kbps**, or **mbps** and the bandwidth value from 1 through 128,000.

```
[edit interfaces]
user@host# set aex aggregated-ether-options minimum-bandwidth bw-unit (bps | gbps | kbps | mbps)
user@host# set aex aggregated-ether-options minimum-bandwidth bw-value value
```

8. Commit the configuration.

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

## SEE ALSO

[link-speed | 1238](#)

[minimum-bandwidth | 1272](#)

## Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces aex]** hierarchy level:

```
[edit interfaces aex]
vlan-tagging;
```

You must also include the **vlan-id** statement:

```
vlan-id number;
```

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*]**

For more information about the **vlan-tagging** and **vlan-id** statements, see [802.1Q VLANs Overview](#).

## SEE ALSO



`vlan-id`  
`vlan-tagging`

## Configuring Untagged Aggregated Ethernet Interfaces

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.

To configure an aggregated Ethernet interface as untagged, remove the `vlan-tagging` statement at the **[edit interfaces aex]** hierarchy level and remove the `vlan-id` statement from 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*]**

**NOTE:** You can configure only one logical interface (unit 0) on the port. The logical unit 0 is used to send and receive LACP or marker protocol data units (PDUs) to and from the individual links.

Table 11 on page 99 lists untagged aggregated Ethernet and LACP support by PIC and router.

**Table 11: Untagged Aggregated Ethernet and LACP Support by PIC and Platform**

| PIC Type                              | M Series         | LACP | T Series | LACP |
|---------------------------------------|------------------|------|----------|------|
| 4-port Fast Ethernet PIC Type 1       | Yes              | Yes  | Yes      | Yes  |
| 1-port Gigabit Ethernet PIC Type 1    | Yes              | Yes  | Yes      | Yes  |
| 2-port Gigabit Ethernet PIC Type 2    | Yes              | Yes  | Yes      | Yes  |
| 4-port Gigabit Ethernet PIC Type 2    | Yes              | Yes  | Yes      | Yes  |
| 1-port 10-Gigabit Ethernet M160       | Yes              | Yes  | NA       | NA   |
| 10-port Gigabit Ethernet PIC Type 3   | Yes (M120, M320) | Yes  | Yes      | Yes  |
| 1-port 10-Gigabit Ethernet PIC Type 3 | N/A              | NA   | Yes      | Yes  |

Table 11: Untagged Aggregated Ethernet and LACP Support by PIC and Platform (*continued*)

| PIC Type                           | M Series | LACP | T Series | LACP |
|------------------------------------|----------|------|----------|------|
| 8-port Gigabit Ethernet PIC Type 3 | Yes      | Yes  | Yes      | Yes  |

The 8-port Fast Ethernet PIC does not support untagged aggregated Ethernet or LACP.

Syslog messages are logged if you try to configure an untagged aggregated Ethernet interface using an unsupported PIC type.

SEE ALSO

[Understanding Ethernet Link Aggregation on ACX Series Routers | 188](#)

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

### IN THIS SECTION

- [Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links Overview | 100](#)
- [Example: Configuring Targeted Distribution for Accurate Policy Enforcement on Logical Interfaces Across Aggregated Ethernet Member Links | 101](#)

This document provides an overview of targeted distribution of static logical interfaces across aggregated Ethernet member links, and an example for configuring targeted distribution.

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

Targeted distribution provides a mechanism to direct traffic through specified links for an aggregated Ethernet bundle, and also assigns roles to member links to handle link failure scenarios.

The targeted distribution of static logical interfaces is also used to accurately enforce egress class-of-service (CoS) profiles. Without this feature, the enforcement of egress class-of-service profiles is distributed among the individual member interface schedulers, shapers, or policers instantiated in each Packet Forwarding Engine that hosts a member link. In the absence of targeted distribution of aggregated Ethernet bundles, traffic destined through a logical interface of a bundle can exit through any of the members based on the hashing algorithm. As a result, it is not possible to determine which link is used to forward traffic. Distributed egress policy enforcement relies on traffic load balancing, which might not be accurate all the time. With targeted distribution, logical interface traffic is directed to a certain member link or a number of member

links. Targeted distribution ensures an accurate policy enforcement that is no longer distributed for a given logical interface.

**NOTE:** Irrespective of the family configured for the logical interface, the targeted distribution feature is applicable to both Layer 3 and Layer 2 interfaces.

You can form distribution lists consisting of member links of the aggregated Ethernet bundle and you can assign roles to these lists, such as primary, backup, and standby. A distribution list specified as primary ensures that traffic is load-balanced among all the 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. 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. 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 primary member links come back online, they continue to carry traffic.

You can configure distribution lists for primary links and backup links. The remaining links are added to a defined standby list. You can make changes to the distribution lists and their roles by configuring them again. When targeted distribution lists are deleted, all links carry traffic. When you commit the configuration, the member links are assigned the specified roles irrespective of whether the links are up or down.

**NOTE:** The feature is supported only on MX Series routers with MPCs, and with the **enhanced-ip** configuration enabled.

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.

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

#### **IN THIS SECTION**

- [Requirements | 102](#)
- [Overview | 102](#)
- [Configuration | 102](#)
- [Verification | 106](#)

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

### 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 Gigabit Ethernet interface, *ge-0/0/3* is specified as distribution list *dl1* and *ge-0/0/4* as *dl2*. The logical interface unit numbers of the aggregated Ethernet bundle ending in an odd number are assigned to the distribution list *dl1* as the primary list, and those ending in an even number are assigned the distribution list *dl2* as the primary list.

To configure targeted distribution you must:

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

### 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 together-options 802.3ad ae10 distribution-list
dl1
user@host# set interfaces ge-0/0/4 apply-groups-except INTF together-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
```

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 together-options 802.3ad ae10 distribution-list
dl1
[edit]
user@host# set interfaces ge-0/0/4 apply-groups-except INTF together-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 provisions.

```

[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;
gether-options {
    802.3ad {
        ae10;
        distribution-list dl1;
    }
}

```

```

user@host# show interfaces ge-0/0/4
apply-groups-except INTF;
gether-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 {
        }
    }
}

```

## Verification

### Verifying 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, *dl1* (**ge-0/0/3**) and those ending in an even number being assigned distribution list, *dl2* (**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: Sendbroadcast-pkt-to-re

```

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

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

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

Protocol inet, MTU: 8978, Generation: 196, Route table: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.2.0.1 , Local: 10.2.0.1, Broadcast: 10.2.0.3, Generation:
152

```

Protocol multiservice, MTU: Unlimited, Generation: 197, Route table: 0  
 Policer: Input: \_\_default\_arp\_policer\_\_

Logical interface ae10.103 (Index 343) (SNMP ifIndex 614) (Generation 152)

Description: matched odd

Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.103 ] Encapsulation: ENET2

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

Bundle:

|         |   |   |     |   |
|---------|---|---|-----|---|
| Input : | 0 | 0 | 0   | 0 |
| Output: | 3 | 0 | 194 | 0 |

Adaptive Statistics:

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

Link:

ge-0/0/3.103

|         |   |   |     |   |
|---------|---|---|-----|---|
| Input : | 0 | 0 | 0   | 0 |
| Output: | 3 | 0 | 194 | 0 |

ge-0/0/4.103

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

| Marker Statistics: | Marker Rx | Resp Tx | Unknown Rx | Illegal Rx |
|--------------------|-----------|---------|------------|------------|
|--------------------|-----------|---------|------------|------------|

|              |   |   |   |   |
|--------------|---|---|---|---|
| ge-0/0/3.103 | 0 | 0 | 0 | 0 |
|--------------|---|---|---|---|

|              |   |   |   |   |
|--------------|---|---|---|---|
| ge-0/0/4.103 | 0 | 0 | 0 | 0 |
|--------------|---|---|---|---|

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

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

Interfaces:

|          |    |
|----------|----|
| ge-0/0/4 | Up |
|----------|----|

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

|        |         |
|--------|---------|
| Backup | Waiting |
|--------|---------|

Interfaces:

|          |    |
|----------|----|
| ge-0/0/3 | Up |
|----------|----|

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

|         |      |
|---------|------|
| Standby | Down |
|---------|------|

Protocol inet, MTU: 8978, Generation: 194, Route table: 0

Flags: Sendbroadcast-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
    Policers: 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 | 1140](#)
- [targeted-distribution | 1381](#)
- [targeted-options | 1382](#)

## Understanding Independent Micro BFD Sessions for LAG

Starting with Junos OS Release 13.3, this feature is supported on the following PIC/FPC types:

- 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
- All MPCs on MX Series with Ethernet MICs
- FPC-PTX-P1-A on PTX5000 with 10-Gigabit Ethernet interfaces
- FPC2-PTX-P1A on PTX5000 with 10-Gigabit Ethernet interfaces in Junos OS Release 14.1 and later
- All FPCs on PTX Series with Ethernet interfaces in Junos OS Release 14.1R3 and later 14.1 releases, and Junos 14.2 and later

**TIP:** See *PTX Series PIC/FPC Compatibility* for a list of PICs that are supported on each PTX Series FPC.

**NOTE:** Micro-BFD configuration with interface addresses is not supported on PTX routers on FPC3 and QFX10000 line of switches.

The Bidirectional Forwarding Detection (BFD) protocol is a simple detection protocol that quickly detects failures in the forwarding paths. A link aggregation group (LAG) combines multiple links between devices that are in point-to-point connections, thereby increasing bandwidth, providing reliability, and allowing load balancing. To run a BFD session on LAG interfaces, 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.

The individual BFD sessions determine the Layer 2 and Layer 3 connectivity of each member link in the LAG. Once a BFD session is established on a particular link, the member links are attached to the LAG and the load balancer either by a static configuration or by the Link Aggregation Control Protocol (LACP). If the member links are attached to the LAG by a static configuration, the device control process acts as the client to the micro BFD session. When member links are attached to the LAG by the LACP, the 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.

**NOTE:**

- 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 accordance with the latest draft for BFD over LAG.
- In Junos OS, MicroBFD control packets are always untagged by default. For L2 aggregated interfaces, the configuration must include vlan-tagging or flexible-vlan-tagging in the Aggregated Ethernet with BFD. Otherwise, the system will throw error while committing the configuration.
- When you enable MicroBFD on an aggregated Ethernet Interface, the aggregated Interface can receive MicroBFD packets. Starting with Junos OS Release 19.3 and later, for MPC10E and MPC11E MPCs, you cannot apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface. For MPC1E through MPC9E, you can apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface only if the aggregated Ethernet Interface is configured as an untagged Interface.

Micro BFD sessions run in the following modes:

- **Distribution Mode**—Micro BFD sessions are distributed by default at Layer 3.
- **Non-Distribution Mode**—You can configure the BFD session to run in this mode by including the **no-delegate-processing** statement under periodic packet management (PPM). In this mode, the packets are being sent or received by the Routing Engine at Layer 2.

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.

**NOTE:** IANA has allocated 6784 as the UDP destination port for micro BFD.

To enable failure detection for LAG networks for aggregated Ethernet interfaces:

- Include the *bfd-liveness-detection* statement in the configuration.
- Specify a hold-down interval value to set the minimum time that the BFD session must remain up before a state change notification is sent to the other members in the LAG network.
- Specify the minimum interval that indicates the time interval for transmitting and receiving data.
- Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases prior to 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.



**NOTE:** On T1600 and T4000 routers, you cannot configure the local aggregated Ethernet Interface address of the remote destination as the neighbor address.



**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 loopback IP address to 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.

**NOTE:** 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.

**NOTE:** 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.

For the IPv6 address family, disable duplicate address detection before configuring this feature with AE interface addresses. To disable duplicate address detection, include the **dad-disable** statement at the **[edit interface aex unit y family inet6]** hierarchy level.

## SEE ALSO

*authentication*

*bfd-liveness-detection*

*detection-time*

*transmit-interval*

## Example: Configuring Independent Micro BFD Sessions for LAG

### IN THIS SECTION

- [Requirements | 116](#)
- [Overview | 116](#)
- [Configuration | 117](#)
- [Verification | 124](#)

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 Trio chipset
- 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**

The example includes two routers that are directly connected. Configure two aggregated Ethernet interfaces, AE0 for IPv4 connectivity and AE1 for IPv6 connectivity. Configure micro BFD session on the AE0 bundle using IPv4 addresses as local and neighbor endpoints on both routers. Configure 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 117](#) shows the sample topology.

Figure 2: Configuring an Independent Micro BFD Session for LAG



### 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, 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

```

## Configuring a Micro BFD Session for Aggregated Ethernet Interfaces

### 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, as per 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 the 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:** Beginning with Junos OS Release 16.1, you can also configure this feature with the AE interface address as the local address in a micro BFD session.

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.

## 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 configuration 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;
  giggerther-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

- [Verifying That the Independent BFD Sessions Are Up | 124](#)
- [Viewing Detailed BFD Events | 126](#)

Confirm that the configuration is working properly.

Verifying 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   Transmit
10.255.106.102          Up        xe-4/0/0     Time    Interval Multiplier
                        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 Time | Transmit 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 Micro BFD field represents the independent micro BFD sessions running on the links in a LAG. The TX interval *item*, RX interval *item* 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 **bfd-liveness-detection** statement.

### Viewing 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]bfdd_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*

*detection-time*

*Configuring Independent Micro BFD Sessions for LAG*

## Deleting an Aggregated Ethernet Interface

There are two approaches to deleting an aggregated Ethernet interface:

- You can delete an aggregated Ethernet interface from the interface configuration. The Junos OS removes the configuration statements related to **aex** and sets this interface to down state.
- You can also permanently remove the aggregated Ethernet interface from the device configuration by deleting it from the device-count on the routing device.

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

This step changes the interface state to down and removing 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
```

#### SEE ALSO

---

[Load Balancing on Aggregated Ethernet Interfaces | 140](#)

[Periodic Packet Management | 183](#)

Release History Table

| Release                | Description                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">19.3</a>   | Starting with Junos OS Release 19.3 and later, for MPC10E and MPC11E MPCs, you cannot apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface. For MPC1E through MPC9E, you can apply firewall filters on the MicroBFD packets received on the aggregated Ethernet Interface only if the aggregated Ethernet Interface is configured as an untagged Interface. |
| <a href="#">19.2R1</a> | Starting with Junos OS Release 19.2R1, you can configure any number of aggregated Ethernet Interfaces in a mixed rate aggregated Ethernet bundle on MX Series routers and PTX Series routers.                                                                                                                                                                                                       |
| <a href="#">18.1</a>   | Starting in Junos OS Release 18.1, Junos OS supports removal of child next hop usage for aggregated Ethernet Interfaces and clients on PTX Series routers with FPC3-PTX-U2 and FPC3-PTX-U3.                                                                                                                                                                                                         |
| <a href="#">16.1</a>   | 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.                                                                                                                                                                                                        |
| <a href="#">16.1</a>   | Beginning with Release 16.1R2, Junos OS checks and validates the configured micro BFD <b>local-address</b> against the interface or loopback IP address before the configuration commit.                                                                                                                                                                                                            |
| <a href="#">15.1F5</a> | In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 480 aggregated interfaces on MX240, MX480, and MX960 routers.                                                                                                                                                                                                                                                                |
| <a href="#">15.1F5</a> | In Junos OS 15.1F5 and 15.1F6 releases, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.                                                                                                                                                                                                                                                                      |
| <a href="#">14.2R4</a> | Starting in Junos OS Release 14.2R4, the timer values that ensure proper link aggregation and STP functions are configured by default if you use Junos Fusion with Junos OS.                                                                                                                                                                                                                        |
| <a href="#">14.2R3</a> | In Junos release 14.2R3 and later, you can configure a maximum of 1000 aggregated interfaces on MX240, MX480, and MX960 routers.                                                                                                                                                                                                                                                                    |
| <a href="#">14.2R3</a> | In Junos release 14.2R3 and later, you can configure a maximum of 800 aggregated interfaces on MX2010 and MX2020 routers.                                                                                                                                                                                                                                                                           |
| <a href="#">14.2R3</a> | If you use Junos Fusion with Junos OS Release 14.2R3, you need to ensure that link aggregation (and STP) work properly by configuring timers for the periodic packet management (PPM) daemons on the aggregation and satellite devices.                                                                                                                                                             |
| <a href="#">14.2R2</a> | In Junos release 14.2R2 and earlier, you can configure a maximum of 480 aggregated interfaces on MX Series routers.                                                                                                                                                                                                                                                                                 |



|        |                                                                                                                                                                                                                                                                                                                                                     |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 14.2R2 | For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. For MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.                                                                                                                                     |
| 14.2   | Starting in Junos OS Release 14.2, you can configure an enhanced link aggregation group (LAG) on MX Series routers.                                                                                                                                                                                                                                 |
| 14.2   | Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.                                                                                                                                                                                                                         |
| 14.2   | Starting with Junos OS Release 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers, thereby enabling you to configure the member links with any combination of rates— 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet— on an aggregated Ethernet bundle. |
| 14.1R1 | Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.                                                                                                                                                                              |
| 14.1   | Starting with Junos OS Release 14.1, you can configure aggregated Ethernet interfaces with LACP on logical systems within an MX Series router.                                                                                                                                                                                                      |
| 14.1   | Starting with Junos OS Release 14.1R1 and 14.2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers.                                                                                                                                                                              |
| 14.1   | Starting with Junos OS Release 14.1, specify the neighbor in a BFD session. In releases prior to 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, you must not configure the LACP system identifier by using the <b>system-id system-id</b> statement at the <b>[edit interfaces aeX aggregated-ether-options lacp]</b> hierarchy level to be all zeros (00:00:00:00:00:00).                                                                                     |
| 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.                                                                                                                                                                                                                               |
| 13.2   | Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP.                                                                                                                                                                         |
| 13.2   | For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 64, numbered from <b>as0</b> through <b>as63</b> .                                                                                                                                                                                                  |
| 13.2   | For SONET/SDH, starting with Junos OS Release 13.2, the maximum number of logical interfaces is 16, numbered from <b>as0</b> through <b>as15</b> .                                                                                                                                                                                                  |

|      |                                                                                                                                                                                                                                                                                                                          |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 13.2 | Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers.                                                                                                                                                                         |
| 13.2 | Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP.                                                                                                                                              |
| 13.2 | Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. |
| 13.2 | Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers.                                                                                                                                                                         |
| 13.2 | Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP.                                                                                                                                              |
| 12.2 | Starting with Junos OS Release 12.2, you can also configure LACP to override the IEEE 802.3ad standard and to allow the standby link always to receive traffic. Overriding the default behavior facilitates subsecond failover.                                                                                          |
| 11.4 | Starting with Junos OS Release 11.4, source address filtering does not work when LACP is enabled.                                                                                                                                                                                                                        |
| 9.3  | Starting with Junos OS Release 9.3, port aggregation selection (discussed previously) is performed for the active link when LACP link protection is enabled.                                                                                                                                                             |

## RELATED DOCUMENTATION

[Performance Monitoring on Aggregated Ethernet Interfaces | 179](#)

[Understanding Ethernet Link Aggregation on ACX Series Routers | 188](#)

## Link Protection of Aggregated Ethernet Interfaces

### IN THIS SECTION

- [Configuring Aggregated Ethernet Link Protection | 133](#)
- [Configuring Aggregated Ethernet Minimum Links | 135](#)
- [Example: Configuring Aggregated Ethernet Link Protection | 137](#)

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 | 134](#)
- [Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces | 134](#)
- [Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link | 135](#)
- [Disabling Link Protection for Aggregated Ethernet Interfaces | 135](#)

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](#) | [1234](#)

[aggregated-ether-options](#) | [1108](#)

### **Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces**

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

To configure a primary link and a backup link:

1. Configure the primary logical interface.

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

2. Configure the backup logical interface.

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

SEE ALSO

| [802.3ad](#) | [992](#)

### ***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 \(revert | switchover\) \(Aggregated Ethernet Link Protection\)](#) | [1482](#)

### ***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 \(revert | switchover\) \(Aggregated Ethernet Link Protection\)](#) | [1482](#)

## **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](#) | 1108

[minimum-links](#) | 1273

## 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 {
  gether-options {
    802.3ad ae0 primary;
  }
}
[edit interfaces]
ge-1/0/1 {
  gether-options {
    802.3ad ae0 backup;
  }
}
```

SEE ALSO

[aggregated-ether-options](#) | 1108

## RELATED DOCUMENTATION

[Load Balancing on Aggregated Ethernet Interfaces](#) | 140

[Performance Monitoring on Aggregated Ethernet Interfaces](#) | 179

## Scheduling on Aggregated Ethernet Interfaces

### IN THIS SECTION

- [Configuring Shared Scheduling on Aggregated Ethernet Interfaces | 138](#)
- [Configuring Scheduler on Aggregated Ethernet Interfaces Without Link Protection | 139](#)

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 | 1108](#)

[link-protection | 1234](#)

*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](#)

For more information on the hierarchical scheduler (CoS), see the *Class of Service User Guide (Routers and EX9200 Switches)*.

## RELATED DOCUMENTATION

[Understanding Ethernet Link Aggregation on ACX Series Routers | 188](#)

[Performance Monitoring on Aggregated Ethernet Interfaces | 179](#)

## Load Balancing on Aggregated Ethernet Interfaces

### IN THIS SECTION

- [Load Balancing and Ethernet Link Aggregation Overview | 141](#)
- [Understanding Aggregated Ethernet Load Balancing | 141](#)
- [Stateful Load Balancing for Aggregated Ethernet Interfaces Using 5-Tuple Data | 144](#)
- [Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces | 147](#)
- [Configuring Adaptive Load Balancing | 148](#)
- [Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers | 149](#)
- [Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs for MX Series Routers | 156](#)
- [Examples: Configuring PIC-Level Symmetrical Hashing for Load Balancing on 802.3ad LAGs on MX Series Routers | 159](#)
- [Example: Configuring Aggregated Ethernet Load Balancing | 161](#)

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 \(Aggregated Ethernet\) | 1553](#)

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

When multiple flows are transmitted out of an aggregated Ethernet (ae) interface, the flows must be distributed across the different member links evenly to enable an effective and optimal load-balancing behavior. To obtain a streamlined and robust method of load-balancing, the member link of the aggregated Ethernet interface bundle that is selected each time for load balancing plays a significant part. In Junos OS releases earlier than Release 13.2R1, on MX Series routers with Trio-based FPCs (MPCs), the selection of a member link of the ae interface bundle or the next-hop (or unilist of next-hops) for equal-cost multipath (ECM) links is performed using a balanced mode next-hop selection methodology and an unbalanced mode of member link or next-hop selection methodology. The balanced mode of link selection uses 'n' bits in a precomputed hash value if it needs to select one of  $2^n$  ( $2$  raised to the power of  $n$ ) next-hop in the unilist. The unbalanced mode of member-link or next-hop selection uses 8 bits in a precomputed hash to select an entry in a selector table, which is randomly done with the member link IDs of the link aggregation group (LAG) or aebundle.

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

The deviation per child link is defined as

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

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 unilist of next-hops in an ECMP link. This method of splitting the load among member links is called *stateful load balancing* and it uses 5-tuple information (source and destination addresses, protocol, source and destination ports). Such a method can be mapped directly to the flows, or to a precompute hash based on certain fields in the flow. As a result, the deviation observed on each child link is reduced.

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

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

To enable per-flow load balancing based on hash values, include the **per-flow** statement at the at the **[edit interfaces 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](#) | 133



## Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces

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

To configure stateful load balancing on **ae** interface bundles:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

SEE ALSO

[Link Protection of Aggregated Ethernet Interfaces](#) | 133

## 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](#) | [1095](#)

## 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](#) | [149](#)
- [Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers](#) | [150](#)
- [Configuring Symmetrical Load Balancing on Trio-Based MPCs](#) | [153](#)
- [Example Configurations](#) | [155](#)

### *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.

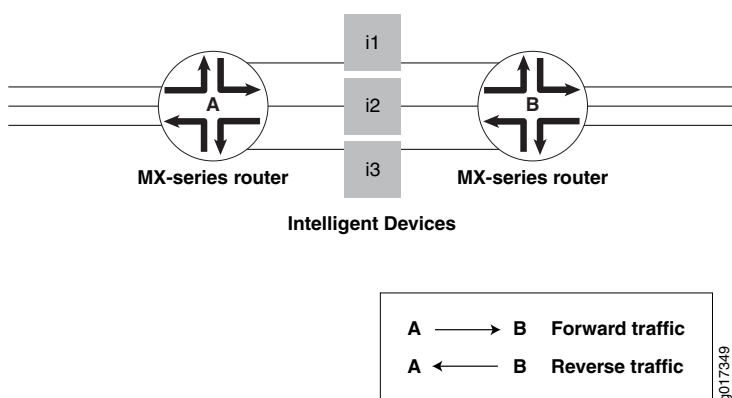
### Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

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 150](#).

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

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

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

**NOTE:**

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

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

For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library*.

### 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 154](#), 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 | 155](#)
- [Example Configurations of Per-Packet-Forwarding-Engine Settings | 155](#)

## 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

For additional information, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS Administration Library*.

## 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 | 159](#)
- [Configuring Symmetrical Hashing for family inet on Both Routers | 160](#)
- [Configuring Symmetrical Hashing for family inet and family multiservice on the Two Routers | 160](#)

**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 | 162](#)

## Example: Configuring Aggregated Ethernet Load Balancing

### IN THIS SECTION

- [Requirements | 162](#)
- [Overview | 162](#)
- [Configuration | 164](#)
- [Verification | 177](#)

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

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

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

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

The aggregated Ethernet load-balancing solutions are:

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

The adaptive load-balancing solution corrects the traffic flow imbalance by modifying the selector entries, and periodically scanning the link utilization on each member link of the AE bundle to detect any deviations. When a deviation is detected, an adjustment event is triggered and fewer flows are mapped to the affected member link. As a result, the offered bandwidth of that member link goes down. This



causes a continuous feedback loop, which over a period of time ensures that the same amount of byte rate is offered to all the member links, thus providing efficient traffic distribution across each member link in the AE bundle.

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

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

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

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

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

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

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

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

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

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


The aggregated Ethernet load-balancing solutions are mutually exclusive. When more than one of the load-balancing solutions is configured, the solution that is configured last overrides the previously configured

one. You can verify the load-balancing solution being implemented by issuing the **show interfaces aeX aggregated-ether-options load-balance** command.

### Topology

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

Figure 5: Aggregated Ethernet Load Balancing



### 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, 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 {
  gicether-options {
    802.3ad ae0;
  }
}

```

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

```

}
ge-2/2/6 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/7 {
  gigether-options {
    802.3ad ae1;
  }
}
ge-2/2/8 {
  gigether-options {
    802.3ad ae1;
  }
}
ae0 {
  aggregated-ether-options {
    load-balance {
      adaptive tolerance 10;
    }
    link-speed 1g;
    lacp {
      active;
    }
  }
  unit 0 {
    family inet {
      address 120.168.104.1/30;
    }
    family iso;
    family mpls;
  }
}
ae1 {
  aggregated-ether-options {
    load-balance {
      adaptive tolerance 10;
    }
    link-speed 1g;
    lacp {
      active;
    }
  }
}

```

```

unit 0 {
    family inet {
        address 120.168.105.1/30;
    }
    family iso;
    family mpls;
}
}
lo0 {
    unit 0 {
        family inet {
            address 120.168.0.4/32;
        }
        family iso {
            address 49.0001.1201.6800.0004.00;
        }
    }
}
}

```

```

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

```

```

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

```

```
interface ge-1/2/1.0;
interface ae0.0;
interface ae1.0;
interface lo0.0;
}
```

Verification

IN THIS SECTION

- [Verifying Adaptive Load Balancing on ae0 | 177](#)

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

| [Configuring Aggregated Ethernet Interfaces and LACP | 57](#)

Release History Table

| Release                | Description                                                                                                                                                                                                                                 |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">14.1</a>   | 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. |
| <a href="#">13.3</a>   | 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.                   |
| <a href="#">13.2R1</a> | 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.                                                |
| <a href="#">10.1</a>   | 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 <b>payload</b> statement.                                            |

RELATED DOCUMENTATION

| [Configuring Aggregated Ethernet Interfaces and LACP | 57](#)

| [Link Protection of Aggregated Ethernet Interfaces | 133](#)



## 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 | 179](#)
- [Guidelines for Configuring Performance Monitoring Functionalities on Aggregated Ethernet Interfaces | 181](#)

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](#) | 689

## 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 | 709](#)

---

[Configuring Ethernet Frame Loss Measurement | 747](#)

---

[Configuring Ethernet Synthetic Loss Measurements | 796](#)

## Release History Table

| Release                | Description                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">16.2</a>   | Starting with Junos OS Release 16.2, connectivity fault management (CFM) is supported when enhanced LAG is enabled.                                                                                                                                                                                                                                                                                            |
| <a href="#">16.1R1</a> | 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.                                                                                                                                       |
| <a href="#">16.1</a>   | 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.                                                                                                                                                                                                                         |
| <a href="#">16.1</a>   | Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the <b>performance-monitoring</b> statement and its substatements at the <b>[edit protocols oam ethernet connectivity-fault-management]</b> 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

[Configuring Aggregated Ethernet Interfaces and LACP | 57](#)

[ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)

## Periodic Packet Management

### IN THIS SECTION

- [Understanding Periodic Packet Management on MX Series Routers | 184](#)
- [Configuring Periodic Packet Management on MX Series Routers | 184](#)

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 | 185](#)
- [Enabling Centralized Periodic Packet Management | 186](#)

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 | 719](#)

---

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

---

[ppm | 1320](#)

## Understanding Ethernet Link Aggregation on ACX Series Routers

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

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

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

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

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

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

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

To configure aggregated Ethernet interface:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

To delete an aggregated Ethernet interface:

1. Delete the aggregated Ethernet configuration.

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

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

2. Delete the interface from the device count.

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

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

## Load Balancing

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

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

ACX Series routers also support load balancing across the member links using Layer 2 source MAC addresses, destination MAC addresses, or both. This can be configured at the **[edit forwarding-options hash-key family multiservice]** hierarchy level. Layer 2 source MAC addresses and destination MAC addresses are used as hash-keys for load balancing.

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

**NOTE:**

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

## LACP Monitoring

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

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

LACP is designed to achieve the following:

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

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

LACP monitoring can be either distributed or centralized. The default is distributed and it can be 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

### IN THIS SECTION

- [Configuring Link Protection for Aggregated Ethernet Interfaces | 193](#)
- [Disabling Link Protection for Aggregated Ethernet Interfaces | 193](#)

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 12 on page 194](#) 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 12: 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 13 on page 195](#) 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 13: 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><br>Layer 4 (UDP/TCP source port andr UDP/TCP destination port) |

RELATED DOCUMENTATION

|                                                                                                |
|------------------------------------------------------------------------------------------------|
| <a href="#">CoS on ACX Series Routers Features Overview</a>                                    |
| <a href="#">Controlling Network Access Using Traffic Policing Overview</a>                     |
| <a href="#">Overview of Firewall Filter Match Conditions and Actions on ACX Series Routers</a> |

# 2

PART

## Gigabit Ethernet Interfaces

---

Configuring Gigabit Ethernet Interfaces | **199**

Configuring Rate Selectability | **313**

Configuring Optical Transport network | **423**

---



# Configuring Gigabit Ethernet Interfaces

## IN THIS CHAPTER

- [Configuring 10-Gigabit Ethernet PICs | 199](#)
- [Configuring 40-Gigabit Ethernet PICs | 238](#)
- [Configuring 100-Gigabit Ethernet MICs/PICs | 243](#)
- [Using Smart SFPs for Transporting Legacy Network Traffic over Packet Switched Networks | 271](#)
- [Configuring Layer 2 Overhead Attribute in Interface Statistics | 283](#)
- [Configuring Gigabit Ethernet Policers | 291](#)
- [Gigabit Ethernet Autonegotiation | 305](#)

## Configuring 10-Gigabit Ethernet PICs

### IN THIS SECTION

- [10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 200](#)
- [12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 204](#)
- [24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview | 206](#)
- [P2-10G-40G-QSFPP PIC Overview | 207](#)
- [Configuring the P2-10G-40G-QSFPP PIC | 218](#)
- [Example: Configuring the P2-10G-40G-QSFPP PIC | 222](#)
- [Framing Overview | 225](#)
- [Understanding WAN Framing | 227](#)
- [Configuring Ethernet Framing | 227](#)
- [Modes of Operation | 228](#)
- [Configuring Mixed-Rate Mode Operation | 229](#)
- [Configuring Line-Rate Mode on PICs Supporting Oversubscription | 230](#)
- [Example: Handling Oversubscription on a 10-Gigabit Ethernet LAN/WAN PIC | 231](#)
- [Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 232](#)

- [Gigabit Ethernet Notification of Link Down Alarm Overview | 236](#)
- [Notification of Link Down for Optics Options Overview | 236](#)
- [Configuring Gigabit Ethernet Notification of Link Down Alarm | 236](#)
- [Configuring Link Down Notification for Optics Options Alarm or Warning | 237](#)

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 14 on page 203](#) 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP).

**Table 14: 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><br>960 filters across multiple logical interfaces |
| MAC policers            | 128 ingress Mac policers<br><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 | 232](#)

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)

[IEEE 802.3ah OAM Link-Fault Management Overview | 655](#)

## 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](#) | 238

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

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

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

- Twenty-four 10-Gigabit Ethernet interfaces in two-to-one oversubscription of traffic in oversubscribed mode or 12 ports in line-rate mode. For more information about oversubscribed mode and line-rate mode, see the [“Configuring Line-Rate Mode on PICs Supporting Oversubscription”](#) on page 230.
- 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 | 291](#)

[Gigabit Ethernet Autonegotiation | 305](#)

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

### IN THIS SECTION

- [Understanding Dual Configuration on P2-10G-40G-QSFPP PIC | 208](#)
- [Understanding Port Group | 209](#)
- [Port Numbering on P2-10G-40G-QSFPP PIC When Port Groups Are Not Configured | 213](#)
- [10-Gigabit Ethernet Mode | 216](#)
- [40-Gigabit Ethernet Mode | 217](#)

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 15 on page 211](#) shows the port numbering in 40-Gigabit Ethernet mode and in 10-Gigabit Ethernet mode at the port group level.

Table 15: Port Number Mapping When Port Groups Are Configured

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



Table 15: Port Number Mapping When Port Groups Are Configured (*continued*)

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

Table 15: 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 |
|-------------------|--------------------------------------------|--------------------------------------------|
| 9(3)              | et-1/1/9                                   | et-1/1/9:0                                 |
|                   |                                            | et-1/1/9:1                                 |
|                   |                                            | et-1/1/9:2                                 |
|                   |                                            | et-1/1/9:3                                 |
|                   | et-1/1/10                                  | et-1/1/10:0                                |
|                   |                                            | et-1/1/10:1                                |
|                   |                                            | et-1/1/10:2                                |
|                   |                                            | et-1/1/10:3                                |
|                   | et-1/1/11                                  | et-1/1/11:0                                |
|                   |                                            | et-1/1/11:1                                |
|                   |                                            | et-1/1/11:2                                |
|                   |                                            | et-1/1/11:3                                |

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

Table 16 on page 213 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 16: Port Number Mapping When Port Groups Are Not Configured

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

Table 16: Port Number Mapping When Port Groups Are Not Configured (*continued*)

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

Table 16: 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               |
|-------------------|--------------------------------------------|----------------------------------------------------------|
| 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     |
| 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

### IN THIS SECTION

- Framing Mode Overview | 216
- Supported Features on LAN PHY and WAN PHY Framing Mode | 217

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](#) | 243

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

### IN THIS SECTION

- [Configuring the PIC in 10-Gigabit Ethernet Mode or in 40-Gigabit Ethernet Mode | 218](#)
- [Configuring the PIC in 10-Gigabit Ethernet Mode to Operate in 40-Gigabit Ethernet Mode | 218](#)
- [Configuring the PIC in 40-Gigabit Ethernet Mode to Operate in 10-Gigabit Ethernet Mode | 219](#)
- [Configuring the PIC at Port Group Level | 220](#)
- [Configuring Framing Mode on P2-10G-40G-QSFPP PIC | 220](#)

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 | 221](#)
- [Configuring LAN PHY Framing Mode in 40-Gigabit Ethernet Mode | 221](#)

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.

SEE ALSO

| [P2-10G-40G-QSFPP PIC Overview | 207](#)

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

### IN THIS SECTION

- [Requirements | 222](#)
- [Overview | 222](#)
- [Configuration | 222](#)

### **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

- [Configuring the P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet Mode | 223](#)
- [Configuring the Framing Mode on an Interface | 223](#)

- [Configuring the Interface Options | 224](#)
- [Verification | 224](#)

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**

#### **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
    Bit errors          0
    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 | 243](#)

### 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 | 243](#)

## 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 | 1194](#)

[Configuring Ethernet Framing | 227](#)

[Framing Overview | 225](#)

*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 | 243](#)

## 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](#) | 1274

## 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 | 228](#)

[mixed-rate-mode | 1274](#)

## 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 | 243](#)

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

[Table 17 on page 231](#) 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 17: 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 ( $C = A \times 2 + B$ ) | Status of Oversubscription and Throughput                                                                 |
|-------------------------------------------------|------------------------------------------------|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| 0                                               | 1                                              | 1                                                          | Oversubscription is not active. Each port will receive 10 Gbps 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). |

Table 17: 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 ( $C = A \times 2 + B$ ) | Status of Oversubscription and Throughput                                                                                                                                                                                                                                                        |
|-------------------------------------------------|------------------------------------------------|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1                                               | 4                                              | 6                                                          | <p>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).</p> <p>For the remaining four ports, oversubscription is not active. Each port will receive 10 Gbps throughput.</p> |
| 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 | 200](#)
[Configuring Line-Rate Mode on PICs Supporting Oversubscription | 230](#)
[Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 232](#)

### 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 233](#).

**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](#) | [1287](#)

## 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](#) | [1113](#)

## Notification of Link Down for Optics Options Overview

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

For information on configuring link down notification, see [“Configuring Link Down Notification for Optics Options Alarm or Warning”](#) on page 237.

SEE ALSO

| [Configuring Link Down Notification for Optics Options Alarm or Warning](#) | [237](#)

## 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 together-options]** hierarchy level:

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

```
    asynchronous-notification;  
  }  
}
```

SEE ALSO

| [asynchronous-notification](#) | [1113](#)

## 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);  
    }  
  }  
}
```

SEE ALSO

| [alarm](#) | [925](#)

| [warning](#) | [983](#)

## Release History Table

| Release                | Description                                                                                                                                                       |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">14.1R2</a> | Starting with Junos OS Release 14.1R2, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC.                                                         |
| <a href="#">12.2</a>   | 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).   |
| <a href="#">12.2</a>   | 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).   |
| <a href="#">12.1R2</a> | 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). |
| <a href="#">12.1R2</a> | 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 | 243](#)

[Configuring Gigabit Ethernet Policers | 291](#)

[Gigabit Ethernet Autonegotiation | 305](#)

## Configuring 40-Gigabit Ethernet PICs

### IN THIS SECTION

- [40-Gigabit Ethernet PIC Overview | 239](#)
- [Configuring 40-Gigabit Ethernet PICs | 241](#)

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

## 40-Gigabit Ethernet PIC Overview

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

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

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

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

- Encapsulation protocols such as:
  - Layer 2 protocols
    - Ethernet CCC, Ethernet TCC, and Ethernet VPLS
    - VLAN CCC
    - Extended VLAN TCC
    - VLAN VPLS
    - Flexible Ethernet service
  - Layer 3 protocols
    - IPv4
    - IPv6
    - MPLS
- CFP Multi-Source Agreement (MSA)-compliant management data input/output (MDIO) control features (transceiver dependent).
- Graceful Routing Engine switchover (GRES) (in all PIC and chassis configurations).
- Interface creation:
  - When the PIC is brought online, the router creates one interface, et-x/y/0, where x represents the FPC slot number and y represents PIC slot number. The physical interface represents internal Ethernet Packet Forwarding Engines.
  - The FPC slot number ranges from 0 through 7 in T640, T1600, and T4000 routers. The PIC slot numbers are 0 and 1.

- Packet Forwarding Engine 0 is the physical interface 0, and Packet Forwarding Engine 1 is the physical interface 1.
- 802.3 link aggregation:
  - The configuration of the 40-Gigabit Ethernet PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
  - An aggregate bundle that consists purely of 40-Gigabit Ethernet PICs supports a maximum of 40-Gigabit Ethernet links depending on the system implementation.

For Junos OS configuration information about this PIC, see [“Configuring 40-Gigabit Ethernet PICs” on page 241](#). 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

|                                             |
|---------------------------------------------|
| <i>T640 Core Router Hardware Guide</i>      |
| <i>T1600 Core Router Hardware Guide</i>     |
| <i>T4000 Core Router Hardware Guide</i>     |
| <i>TX Matrix Plus Router Hardware Guide</i> |
| <i>T640 PICs Supported</i>                  |
| <i>T1600 PICs Supported</i>                 |
| <i>T4000 PICs Supported</i>                 |

## 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” on page 64](#).
- 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” on page 64](#) and [“10-port 10-Gigabit Ethernet LAN/WAN PIC Overview” on page 200](#).

4. Configure the Packet Forwarding Engine features.

The 40-Gigabit Ethernet PIC with CFP supports all classification, firewall filters, queuing model, and rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see [“Configuring Gigabit Ethernet Policers” on page 293](#), [“Configuring Gigabit Ethernet Policers” on page 293](#), and *Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*.

## SEE ALSO

[Configuring Gigabit Ethernet Policers | 293](#)

[Configuring Gigabit Ethernet Policers | 293](#)

*Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview*

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 291](#)

[Gigabit Ethernet Autonegotiation | 305](#)

## Configuring 100-Gigabit Ethernet MICs/PICs

### IN THIS SECTION

- 100-Gigabit Ethernet Interfaces Overview | 243
- MPC3E MIC Overview | 246
- 100-Gigabit Ethernet Type 4 PIC with CFP Overview | 247
- Configuring 100-Gigabit Ethernet Type 4 PIC With CFP | 251
- Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 256
- 100-Gigabit Ethernet Type 5 PIC with CFP Overview | 259
- 100-Gigabit Ethernet Interfaces Interoperability | 261
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | 263
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 264
- Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | 266
- Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 267

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 | 243
- PTX Series 100-Gigabit Ethernet Interfaces | 244
- T Series 100-Gigabit Ethernet Interfaces | 245

#### *MX Series 100-Gigabit Ethernet Interfaces*

[Table 18 on page 244](#) lists the 100-Gigabit Ethernet interfaces supported by MX Series routers.



Table 18: 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 | 100-Gigabit Ethernet MIC with CFP<br><br><a href="#">“MPC3E MIC Overview” on page 246</a> |
| 100-Gigabit Ethernet MIC with CXP           | MIC3-3D-1X100GE-CXP | MX240<br>MX480<br>MX960<br>MX2010<br>MX2020 | 100-Gigabit Ethernet MIC with CXP<br><br><a href="#">“MPC3E MIC Overview” on page 246</a> |
| 100-Gigabit Ethernet ports on the MPC4E     | MPC4E-3D-2CGE-8XGE  | MX240<br>MX480<br>MX960<br>MX2010<br>MX2020 | MPC4E on MX Series Routers Overview<br><br>2x100GE + 8x10GE MPC4E                         |
| 100-Gigabit Ethernet MIC with CFP2          | MIC6-100G-CFP2      | MX2010<br>MX2020                            | 100-Gigabit Ethernet MIC with CFP2                                                        |
| 100-Gigabit Ethernet MIC with CXP (4 Ports) | MIC6-100G-CXP       | MX2010<br>MX2020                            | 100-Gigabit Ethernet MIC with CXP (4 Ports)                                               |

**PTX Series 100-Gigabit Ethernet Interfaces**

[Table 19 on page 244](#) lists the 100-Gigabit Ethernet interfaces supported by PTX Series routers.

Table 19: PTX Series 100-Gigabit Ethernet Interfaces

| PIC                               | Model Number       | Routers Supported | For More Information                           |
|-----------------------------------|--------------------|-------------------|------------------------------------------------|
| 100-Gigabit Ethernet PIC with CFP | P1-PTX-2-100GE-CFP | PTX5000           | 100-Gigabit Ethernet PIC with CFP (PTX Series) |

Table 19: PTX Series 100-Gigabit Ethernet Interfaces (*continued*)

| PIC                                | Model Number      | Routers Supported  | For More Information                                                                                                                                                                                                   |
|------------------------------------|-------------------|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 100-Gigabit Ethernet PIC with CFP2 | P2-100GE-CFP2     | PTX5000            | <i>100-Gigabit Ethernet PIC with CFP2 (PTX Series)</i>                                                                                                                                                                 |
| 100-Gigabit Ethernet OTN PIC       | P2-100GE-OTN      | PTX5000            | <i>100-Gigabit Ethernet OTN PIC with CFP2 (PTX Series)</i><br><br><a href="#">“Understanding the P2-100GE-OTN PIC” on page 433</a><br><br><a href="#">“Configuring OTN Interfaces on P2-100GE-OTN PIC” on page 516</a> |
| 100-Gigabit DWDM OTN PIC           | P1-PTX-2-100G-WDM | PTX5000<br>PTX3000 | <i>100-Gigabit DWDM OTN PIC (PTX Series)</i>                                                                                                                                                                           |

**T Series 100-Gigabit Ethernet Interfaces**

Table 20 on page 245 lists the 100-Gigabit Ethernet interfaces supported by T Series routers.

Table 20: 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    | <i>100-Gigabit Ethernet PIC with CFP (T1600 Router)</i><br><br><i>100-Gigabit Ethernet PIC with CFP (T4000 Router)</i><br><br><a href="#">“100-Gigabit Ethernet Type 4 PIC with CFP Overview” on page 247</a><br><br><a href="#">“Configuring 100-Gigabit Ethernet Type 4 PIC With CFP” on page 251</a> |

Table 20: T Series 100-Gigabit Ethernet Interfaces (*continued*)

| PIC                                        | Model Number | Routers Supported | For More Information                                                                                                                    |
|--------------------------------------------|--------------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| 100-Gigabit Ethernet PIC with CFP (Type 5) | PF-1CGE-CFP  | T4000             | 100-Gigabit Ethernet PIC with CFP (T4000 Router)<br><br><a href="#">“100-Gigabit Ethernet Type 5 PIC with CFP Overview” on page 259</a> |

## SEE ALSO

---

*MICs Supported by MX Series Routers*

---

*MPCs Supported by MX Series Routers*

---

*PICs Supported on the PTX Series*

---

*T1600 PICs Supported*

---

*T4000 PICs Supported*

**MPC3E MIC Overview**

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

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

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

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

The MPC3E also supports these legacy MICs:

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

The 100-Gigabit Ethernet CFP MIC supports the IEEE standards—compliant 100BASE-LR4 interface, using the 100G CFP optical transceiver modules for connectivity. The 100-Gigabit Ethernet CXP MIC supports

the 100BASE-SR10 interface, using 100-Gigabit CXP optical transceiver modules for connectivity. The 2-port 40-Gigabit Ethernet 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

[MPC3E on MX Series Routers Overview](#)

[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)*

---

*100-Gigabit Ethernet PIC with CFP (T4000 Router)*

## Configuring 100-Gigabit Ethernet Type 4 PIC With CFP

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

- Flexible Ethernet services encapsulation
- Source address MAC filtering
- Destination address MAC filtering
- MAC accounting in RX
- Channels defined by two stacked VLAN tags
- Channels defined by flex-vlan-tagging
- IP service for stacked VLAN tags
- Layer 2 rewrite

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

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

### NOTE:

- For the 100-Gigabit Ethernet Type 4 PIC with CFP, only the 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” on page 293](#), [“Configuring Gigabit Ethernet Policers” on page 293](#), 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 | 293](#)

[Configuring Gigabit Ethernet Policers | 293](#)

[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 PFEO.

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

```

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

```

## SEE ALSO

[forwarding-mode \(100-Gigabit Ethernet\) | 1189](#)

[sa-multicast \(100-Gigabit Ethernet\) | 1344](#)

[vlan-rule \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1424](#)

[vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1425](#)

## 100-Gigabit Ethernet Type 5 PIC with CFP Overview

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

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

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

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

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



- The following Layer 3 protocols are also supported:
  - IPv4
  - IPv6
  - MPLS
- CFP Multi-Source Agreement (MSA) compliant Management Data Input/Output (MDIO) control features (transceiver dependent).
- 802.3 link aggregation:
  - The configuration of the 100-Gigabit Ethernet Type 5 PIC with CFP complies with that of the existing 1-Gigabit or 10-Gigabit Ethernet PIC and aggregated Ethernet interfaces.
- Interoperability mode—Interoperability with the 100-Gigabit Ethernet Type 4 PIC with CFP through configuration in **sa-multicast** forwarding mode.
- Juniper Networks enterprise-specific Ethernet Media Access Control (MAC) MIB
- The 100-Gigabit Ethernet Type 5 PIC with CFP supports all Gigabit Ethernet PIC classification, firewall filters, queuing model, and Layer 2 rewrite functionality features of the Gigabit Ethernet PICs. To configure these parameters, see [“Configuring Gigabit Ethernet Policers” on page 293](#), [“Configuring Gigabit Ethernet Policers” on page 293](#), 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 21 on page 260](#) lists the capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP.

**Table 21: Capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP**

| Capability                         | Support |
|------------------------------------|---------|
| Maximum logical interfaces per PIC | 4093    |

**Table 21: Capabilities of 100-Gigabit Ethernet Type 5 PIC with CFP (*continued*)**

| Capability                          | Support                                                        |
|-------------------------------------|----------------------------------------------------------------|
| Maximum logical interfaces per port | For IPv4 the limit is 4093.<br><br>For IPv6 the limit is 1022. |

SEE ALSO

|                                                                   |
|-------------------------------------------------------------------|
| <a href="#">Configuring Gigabit Ethernet Policers   293</a>       |
| <a href="#">Configuring Gigabit Ethernet Policers   293</a>       |
| <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview</i> |

**100-Gigabit Ethernet Interfaces Interoperability**

**IN THIS SECTION**

- [Interoperability of the MIC-3D-1X100GE-CFP MIC with PICs on Other Routers | 261](#)
- [Interoperability of the MPC4E-3D-2CGE-8XGE MPC with PICs on Other Routers | 262](#)
- [Interoperability of the P1-PTX-2-100GE-CFP PIC with PICs on Other Routers | 262](#)
- [Interoperability of the PD-1CE-CFP-FPC4 PIC with PICs or MICs on Other Routers | 262](#)

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 22 on page 261](#) lists the Interoperability with the 100-Gigabit Ethernet MIC with CFP.

**Table 22: 100-Gigabit Ethernet MIC with CFP (MIC3-3D-1X100GE-CFP) Interoperability**

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

**Table 22: 100-Gigabit Ethernet MIC with CFP (MIC3-3D-1X100GE-CFP) Interoperability (continued)**

|          |                                                               |                                                                                                                                              |
|----------|---------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| T Series | 100-Gigabit Ethernet PIC with CFP (Type 4) (PD- 1CE-CFP-FPC4) | <i>Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode</i> |
|----------|---------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|

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

Table 23 on page 262 lists the Interoperability with the MPC4E.

**Table 23: MPC4E Interoperability**

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

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

Table 24 on page 262 lists the Interoperability with 100-Gigabit Ethernet PIC with CFP (Type 5).

**Table 24: 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) | <a href="#">“Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP” on page 266</a><br><br><a href="#">“Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4” on page 267</a> |

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

Table 25 on page 262 lists the 100-Gigabit Ethernet PIC with CFP (Type 4).

**Table 25: 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) | <a href="#">“Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4” on page 264</a><br><br><a href="#">forwarding-mode</a><br><br><a href="#">sa-multicast</a> |

**Table 25: 100-Gigabit Ethernet PIC with CFP (Type 4) PD-1CE-CFP-FPC4 Interoperability (continued)**

|            |                                                                  |                                                                                                                                                                                                                                                                                |
|------------|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MX Series  | 100-Gigabit Ethernet MIC with CFP (MIC3-3D- 1X100GE-CFP)         | <i>Configuring 100-Gigabit Ethernet MICs to Interoperate with Type 4 100-Gigabit Ethernet PICs (PD-1CE-CFP-FPC4) Using SA Multicast Mode</i>                                                                                                                                   |
|            | 100-Gigabit Ethernet ports on the MPC4E                          | <i>Configuring MPC4E (MPC4E-3D-2CGE-8XGE) to Interoperate with 100-Gigabit Ethernet PICs on Type 4 FPC Using SA Multicast Mode</i>                                                                                                                                             |
| PTX Series | 100-Gigabit Ethernet PIC with CFP (Type 5) (P1- PTX-2-100GE-CFP) | <a href="#">“Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP” on page 266</a><br><br><a href="#">“Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4” on page 267</a> |

SEE ALSO

[Periodic Packet Management | 183](#)
[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](#) | [1189](#)

[sa-multicast](#) | [1344](#)

*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](#) | [264](#)

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)

## SEE ALSO

---

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](#) | [1346](#)

## 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 | 267](#)
- [Configuring Two 50-Gigabit Ethernet Physical Interfaces on the 100-Gigabit Ethernet PIC PD-1CE-CFP-FPC4 as One Aggregated Ethernet Interface | 269](#)

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 gigheter-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 gigheter-options 802.3ad ae0
user@host# set et-5/0/0:1 gigheter-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 | 77](#)

[802.3ad | 992](#)

SEE ALSO

[sa-multicast | 1346](#)

Release History Table

| Release              | Description                                                                                                                |
|----------------------|----------------------------------------------------------------------------------------------------------------------------|
| <a href="#">13.2</a> | Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on 100-Gigabit Ethernet PIC. |

RELATED DOCUMENTATION

|                                                             |
|-------------------------------------------------------------|
| <a href="#">Configuring Gigabit Ethernet Policers   291</a> |
| <a href="#">Gigabit Ethernet Autonegotiation   305</a>      |

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

IN THIS SECTION

- [Transporting Legacy Traffic over Packet Switched Networks | 272](#)
- [Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH Traffic | 274](#)
- [Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy SDH Traffic | 278](#)

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

## Transporting Legacy Traffic over Packet Switched Networks

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, and MPC3 line 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 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, and MPC3 line cards support the smart SFP transceivers.

**NOTE:** The MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE) line card supports the STM16 smart SFP transceiver.

The smart SFP transceivers encapsulate traffic on SDH interfaces on the WAN side as Ethernet frames on the system side. The encapsulated traffic from the SFP transceivers is sent over the PSNs provisioned across the network. You can encapsulate the SDH interfaces using MEF8 framing. You can also configure single VLAN tagging within a default emulation circuit identifier (ECID).

You can encapsulate STM traffic as Ethernet frames using MEF8 framing. Only single VLAN tagging is supported. STM encapsulation uses the Transparent SONET/SDH over Packet (TSoP) standard.

At the local end, the smart SFP transceiver slices the TDM data stream, encapsulates the Ethernet frames and pushes it onto the PSN. The smart SFPs are always paired on the other end of the emulated circuit, and are preconfigured to be in the same multicast MAC address group. At the far end, the smart SFP transceiver decapsulates the Ethernet frames, re-builds the TDM data stream and forwards it onto the local TDM interface.

### ***Benefits of Smart SFP Transceivers***

- Lower operational costs—Smart SFP transceivers enable easy and simplified migration and upgrades from legacy networks to PSNs.
- Operational simplicity and flexibility—You don't need to configure individual TDM interfaces over packet connections. You can deploy additional equipment only if you need it.
- Space saving. Does not require additional rack space.

- Low carbon footprint. Lower power consumption and existing equipment such as TDM access nodes are still in use after migration. Reduced electronic waste.
- Definitive migration to only-Ethernet based network equipment, removing the need to have dedicated TDM network interface cards for the termination of the TDM lines.

### Example: Configuring the Smart-SFPs on MX Series Routers for transporting legacy PDH Traffic

#### IN THIS SECTION

- [Requirements for Configuration of the Smart-SFPs on MX Series Routers | 274](#)
- [Overview | 274](#)
- [Configuring the DS3 Smart SFP | 274](#)
- [Verification | 277](#)

#### *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*

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

In this example, you configure the Smart SFP to transport PDH traffic over PSN networks. To configure the Smart SFP, perform the following tasks:

1. In Configuration mode, create a valid Interface to enable the Smart SFP to communicate with the Junos OS. Configuring VLAN tagging creates a control interface.

```
[edit]
user@host #set interfaces ge-4/0/0 unit 0
user@host #set interfaces ge-4/0/0 vlan-tagging
```

2. Specify the type of Smart SFP to be configured on the interface. In this example, we are configuring a DS3 Smart SFP.

```
[edit]
user@host #set interfaces ge-4/0/0 tdm-options sfp-type DS3
```

3. (Optional) Configure the destination MAC address on the local end smart SFP using the **dmac-address** statement at the **[edit interfaces ge-4/0/0 tdm-options]** hierarchy level to encapsulate the MAC address of the far end smart SFP. To enable MAC address validation or checking of the destination MAC address at the far end smart SFP, use the **ces-psn-port-dmac-check-enable** statement. If the MAC address of the packet does not match, the packet is discarded.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-channel dmac-address dmac-address
```

4. (Optional) Configure the encapsulation mode (MEF8 or MPLS) for further network processing. The default encapsulation mode for DS3 Smart SFP is MEF8.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-channel mode mode
```

5. (Optional) Configure single or dual VLAN tagging on the encapsulated packets. DS3 Smart SFP supports both single and dual VLAN tagging. If you want to configure single VLAN tagging, use the **vlan-id-1** statement and specify the VLAN ID. If you want to configure dual VLAN tagging, use **vlan-id-1** and **vlan-id-2** statements to configure the inner and outer VLAN IDs. Possible values for the VLAN ID: 0 through 4094.

#### Single VLAN tagging

```
[edit interfaces ge-4/0/0 tdm-options]
```



```
user@host # set ces-psn-channel vlan-id-1 vlan-id
```

### Dual VLAN tagging

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-channel vlan-id-1 vlan-id
user@host # set ces-psn-channel vlan-id-2 vlan-id
```

6. (Optional) Configure the emulation circuit ID for encapsulation and decapsulation. If you do not specify an emulation circuit ID, the default value is 0. Possible values for the encapsulation and decapsulation ID: 0 through 1048575.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set iwf-params encap-ecid encap-ecid
user@host # set iwf-params decap-ecid decap-ecid
```

7. (Optional) Specify if you require checking of the destination MAC address of the incoming packets on the receiving SFP at the `[edit interfaces ge-4/0/0 tdm-options]` hierarchy. If you have configured the destination MAC address using the `dmac-address` option, use this option to verify the MAC address on the receiving SFP. If you have enabled MAC address verification and the MAC address does not match, the packet is discarded by the smart SFP.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set ces-psn-port-dmac-check-enable
```

8. (Optional) Enable looping back of the input path of TDM traffic on the SFP TDM port. The input path refers to the traffic from the TDM side that is looped back.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set tdm-in-loop
```

9. (Optional) Enable looping back of the output path of TDM traffic on the SFP TDM port. The output path refers to the traffic from the Ethernet side that is looped back.

```
[edit interfaces ge-4/0/0 tdm-options]
user@host # set tdm-out-loop
```

# Verification

## IN THIS SECTION

- [Verifying the DS3 Smart SFP Statistics on the Interface | 277](#)

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 | 279](#)
- [Overview | 279](#)
- [Configuring the STM1 Smart SFP | 279](#)
- [Verification | 281](#)

### 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

#### 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 | 281](#)

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 | 283](#)
- [Configuring Layer 2 Overhead Accounting in Interface Statistics | 286](#)
- [Verifying the Accounting of Layer 2 Overhead in Interface Statistics | 287](#)

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

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 26 on page 284](#) 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 26: 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](#) | 1091

## Configuring Layer 2 Overhead Accounting in Interface Statistics

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 | 286](#)

#### *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](#) | [1091](#)

## Verifying the Accounting of Layer 2 Overhead in Interface Statistics

### 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 | 1656](#)
- show interfaces statistics*

Release History Table

| Release              | Description                                                                                                                                                                          |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">13.3</a> | Starting with Junos OS Release 13.3, <b>account-layer2-overhead</b> is not supported on MX Series routers with MPC3E (on both PIC and logical interface levels).                     |
| <a href="#">13.2</a> | 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 | 3](#)
- [Initial Configuration of Ethernet Interfaces | 4](#)
- [Understanding Optical Transport Network \(OTN\) | 423](#)

## Configuring Gigabit Ethernet Policers

### IN THIS SECTION

- [Capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs | 291](#)
- [Configuring Gigabit Ethernet Policers | 293](#)
- [Configuring Gigabit Ethernet Two-Color and Tricolor Policers | 300](#)

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 27 on page 292](#) 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 27: 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                     |
| 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                    |

Table 27: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs (continued)

| Capability              | Gigabit Ethernet IQ (SFP) | Gigabit Ethernet (SFP) |
|-------------------------|---------------------------|------------------------|
| <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 *Class of Service User Guide (Routers and EX9200 Switches)*.

SEE ALSO

|                                                         |
|---------------------------------------------------------|
| <a href="#">Configuring MAC Address Accounting   24</a> |
| <a href="#">Configuring a Policer Overhead</a>          |

Configuring Gigabit Ethernet Policers

IN THIS SECTION

- [Overview | 294](#)
- [Configuring a Policer | 294](#)
- [Specifying an Input Priority Map | 295](#)
- [Specifying an Output Priority Map | 295](#)
- [Applying a Policer | 296](#)
- [Configuring MAC Address Filtering | 298](#)
- [Example: Configuring Gigabit Ethernet Policers | 298](#)

## 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 gigether-options ethernet-switch-profile]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
ethernet-policer-profile {
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
}
```

In the Ethernet policer profile, the aggregate-priority policer is mandatory; the premium-priority policer is optional.

For aggregate and premium policers, you specify the bandwidth limit in bits per second. You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). There is no absolute minimum value for bandwidth limit, but any value below 61,040 bps will result in an effective rate of 30,520 bps. The maximum bandwidth limit is 4.29 Gbps.

The maximum burst size controls the amount of traffic bursting allowed. To determine the burst-size limit, you can multiply the bandwidth of the interface on which you are applying the filter by the amount of time you allow a burst of traffic at that bandwidth to occur:

$$\text{burst size} = \text{bandwidth} \times \text{allowable time for burst traffic}$$

If you do not know the interface bandwidth, you can multiply the maximum MTU of the traffic on the interface by 10 to obtain a value. For example, the burst size for an MTU of 4700 would be 47,000 bytes. The burst size should be at least 10 interface MTUs. The maximum value for the burst-size limit is 100 MB.

### **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 **ieee802.1p premium** statement at the **[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]** hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]
ieee802.1p premium [ values ];
```

The priority values can be from 0 through 7. The remaining traffic is classified as nonpremium (or aggregate). For a configuration example, see [“Example: Configuring Gigabit Ethernet Policers” on page 298](#).

**NOTE:** On IQ2 and IQ2-E interfaces and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

### **Specifying an Output Priority Map**

An output priority map identifies egress traffic with specified queue classification and packet loss priority (PLP), and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an output priority map by including the **classifier** statement at the [edit interfaces *interface-name* gigether-options ethernet-policer-profile **output-priority-map**] hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile output-priority-map]
classifier {
  premium {
    forwarding-class class-name {
      loss-priority (high | low);
    }
  }
}
```

You can define a forwarding class, or you can use a predefined forwarding class. [Table 28 on page 296](#) shows the predefined forwarding classes and their associated queue assignments.

Table 28: 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 *Class of Service User Guide (Routers and EX9200 Switches)*. For a configuration example, see [“Example: Configuring Gigabit Ethernet Policers” on page 298](#).

**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” on page 296](#).

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

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

For more information about the **source-filtering** statement, see [“Configuring MAC Address Filtering for Ethernet Interfaces” on page 21](#).

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* *gether-options*]** hierarchy level:

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

### Example: Configuring Gigabit Ethernet Policers

#### IN THIS SECTION

- [Example | 299](#)
- [Example Configuration | 299](#)

### 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 {
  gige-ether-options {
    ether-switch-profile {
      ether-policer-profile {
        input-priority-map {
          ieee-802.1p {
            premium [ 2 3 ];
          }
        }
      }
      output-priority-map {
        classifier {
          premium {
            forwarding-class best-effort {
              loss-priority low;
            }
            forwarding-class expedited-forwarding {
              loss-priority low;
            }
            forwarding-class assured-forwarding {
              loss-priority high;
            }
            forwarding-class network-control {
              loss-priority high;
            }
          }
        }
      }
    }
  }
}
```



```

    policer policer-1 {
      premium {
        bandwidth-limit 100m;
        burst-size-limit 3k;
      }
      aggregate {
        bandwidth-limit 200m;
        burst-size-limit 3k;
      }
    }
  }
}
unit 0 {
  accept-source-mac {
    mac-address 00:01:02:03:04:05 {
      policer {
        input policer-1;
        output policer-1;
      }
    }
  }
}
}

```

## RELATED DOCUMENTATION

[Configuring MAC Address Accounting | 24](#)

*Configuring a Policer Overhead*

## Configuring Gigabit Ethernet Two-Color and Tricolor Policers

### IN THIS SECTION

- [Overview | 301](#)
- [Configuring a Policer | 302](#)
- [Applying a Policer | 303](#)
- [Example: Configuring and Applying a Policer | 303](#)

## Overview

For Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers, you can configure two-color and tricolor marking policers and apply them to logical interfaces to prevent traffic on the interface from consuming bandwidth inappropriately.

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

Juniper Networks router architectures support three types of policer:

- **Two-color policer**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit in some way, or simply discard them. A policer is most useful for metering traffic at the port (physical interface) level.
- **Single-rate tricolor marking (single-rate TCM)**—A single-rate tricolor marking policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).

Starting in Junos OS Release 13.1, traffic is classified into three categories: Green, Red, and Yellow. Following list describes the categories:

- **Green**—Burst size of the packets that arrive is less than the sum of the configured CIR and CBS.
- **Red**—Burst size of the packets that arrive is greater than the sum of the configured CIR and EBS.
- **Yellow**—Burst size of the packets that arrive is greater than the CBS but less than the EBS.

Single-rate TCM is most useful when a service is structured according to packet length and not peak arrival rate.

- **Two-rate Tricolor Marking (two-rate TCM)**—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and peak information rate (PIR), along with their associated burst sizes, the CBS and EBS.

Traffic is classified into the following three categories:

- **Green**—Burst size of the packets that arrive is less than the sum of the configured CIR and CBS.
- **Red**—Burst size of the packets that arrive is greater than the sum of the configured PIR and EBS.

- Yellow—Traffic does not belong to either the green or the red category.

Two-rate TCM is most useful when a service is structured according to arrival rates and not necessarily packet length.

**NOTE:** Unlike policing (described in [“Configuring Gigabit Ethernet Policers” on page 293](#)), 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 *Class of Service User Guide (Routers and EX9200 Switches)*.

### 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

|                                                         |
|---------------------------------------------------------|
| <a href="#">Configuring MAC Address Accounting   24</a> |
| <i>Configuring a Policer Overhead</i>                   |

Release History Table

| Release              | Description                                                                                             |
|----------------------|---------------------------------------------------------------------------------------------------------|
| <a href="#">13.1</a> | Starting in Junos OS Release 13.1, traffic is classified into three categories: Green, Red, and Yellow. |

## RELATED DOCUMENTATION

[Configuring 10-Gigabit Ethernet PICs | 199](#)[Configuring 40-Gigabit Ethernet PICs | 238](#)

## Gigabit Ethernet Autonegotiation

## IN THIS SECTION

- [Gigabit Ethernet Autonegotiation Overview | 305](#)
- [Configuring Gigabit Ethernet Autonegotiation | 306](#)

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 | 243](#)

## Configuring Gigabit Ethernet Autonegotiation

### IN THIS SECTION

- [Configuring Gigabit Ethernet Autonegotiation with Remote Fault | 306](#)
- [Configuring Flow Control | 306](#)
- [Configuring Autonegotiation Speed on MX Series Routers | 306](#)
- [Displaying Autonegotiation Status | 307](#)

### *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 gether-options]** hierarchy level.

```
[edit interfaces ge-fpc/pic/port gether-options]
(auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online | local-interface-offline>
```

### *Configuring Flow Control*

To enable flow control, include the **flow-control** statement at the **[edit interfaces ge-fpc/pic>/port gether-options]** hierarchy level. For more information, see [“Configuring Flow Control” on page 19](#).

### *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 29 on page 307](#) and [Table 30 on page 310](#) 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 29: Mode and Autonegotiation Status (Local)**

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



Table 29: Mode and Autonegotiation Status (Local) (continued)

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

Table 29: Mode and Autonegotiation Status (Local) (continued)

| Transmit | Receive | Mode                 | LED   | Link | Autonegotiation Status |
|----------|---------|----------------------|-------|------|------------------------|
| ON       | ON      | Explicit+RFI-Offline | Green | UP   | No-autonegotiation     |
| ON       | OFF     | Explicit+RFI-Offline | Red   | DOWN |                        |
| OFF      | ON      | Explicit+RFI-Offline | Green | UP   |                        |
| OFF      | OFF     | Explicit+RFI-Offline | Red   | DOWN |                        |
| ON       | ON      | Explicit+RFI-Offline | Red   | DOWN |                        |
| ON       | ON      | Explicit+RFI-Offline | Red   | DOWN | Complete               |
| ON       | OFF     | Explicit+RFI-Offline | Red   | DOWN |                        |
| OFF      | ON      | Explicit+RFI-Online  | Red   | DOWN |                        |
| 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 29: Mode and Autonegotiation Status (Local) (continued)

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

Table 30: Mode and Autonegotiation Status (Remote)

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

Table 30: Mode and Autonegotiation Status (Remote) (*continued*)

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

SEE ALSO

[Configuring 10-Gigabit Ethernet PICs | 199](#)
[Configuring Layer 2 Overhead Attribute in Interface Statistics | 283](#)
**Release History Table**

| Release              | Description                                                                                                                                                                                                                                                                                                     |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">14.2</a> | Starting from Junos OS Release 14.2, the <b>auto-10m-100m</b> option allows the fixed tri-speed port to auto negotiate with ports limited by <b>100m</b> or <b>10m</b> maximum speed.                                                                                                                           |
| <a href="#">14.2</a> | 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 <b>auto-10m-100m</b> 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 | 238](#)

[Configuring 100-Gigabit Ethernet MICs/PICs | 243](#)

# Configuring Rate Selectability

## IN THIS CHAPTER

- [Introduction to Rate Selectability | 313](#)
- [Interface Naming Conventions for Rate Selectability | 345](#)
- [Preventing Oversubscription Using Active Physical Ports | 364](#)
- [Configuring Rate Selectability | 379](#)

## Introduction to Rate Selectability

### IN THIS SECTION

- [Understanding Rate Selectability | 314](#)
- [Guidelines for Configuring Rate Selectability | 318](#)
- [MX10003 MPC Rate-Selectability Overview | 320](#)
- [MX204 Router Rate-Selectability Overview | 325](#)
- [MPC10E-15C-MRATE Rate-Selectability Overview | 332](#)
- [MPC10E-10C-MRATE Rate-Selectability Overview | 336](#)
- [PTX10003 Router Rate-Selectability Overview | 340](#)
- [MX2K-MPC11E Rate Selectability Overview | 342](#)

Use this topic to understand more about rate selectability or multi-rate ports in a network device or in a network component such as a line card.

## Understanding Rate Selectability

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 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. For instance, if a Modular Port Concentrator (MPC) supports multiple speeds, it is known as a rate-selectable MPC. If a Modular Interface Card (MIC) supports multiple speeds, it is known as a rate-selectable MIC. The term *multi-rate* in the name of a component also indicates support for more than one speed.

Rate selectability enables you to configure the port speed either at the port level or at the PIC or MIC level. To configure all ports to operate at the same speed, you configure rate selectability at the MIC or PIC level. 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.

When you configure rate selectability at the MIC or PIC level, all the ports of the MIC or PIC that support the configured speed operate at that speed. To prevent switch fabric interface oversubscription—for example, with the Switch Fabric Board SFB or SFB2— and to ensure a guaranteed bandwidth, you can specify the number of active ports that operate at the configured speed. For instance, on a router with SFB, if you want only eight ports of the MIC to operate at 40 Gbps, you can configure the MIC to operate at 40 Gbps and enable the eight ports that you want to operate at that speed. The remaining ports of the MIC are automatically disabled. For example, 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. Configuring rate selectability at the MIC level helps you configure the operating speed of the MIC easily.

**NOTE:** The total port speed of the MIC cannot exceed the forwarding capacity of the Packet Forwarding Engine.

Configuring rate selectability at the port level provides you the flexibility of operating the ports of the MIC at different supported speeds. 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.

**NOTE:** When you configure rate selectability at the port level, ensure that you plug in transceivers to the ports according to the speeds that you configure. For instance, use 4x duplex LC breakout transceivers to configure 10-Gigabit Ethernet interfaces, fiber-optic 40-gigabit QSFP+ transceivers to configure 40-Gigabit Ethernet interfaces, and fiber-optic 100-gigabit QSFP28 transceivers to configure 100-Gigabit Ethernet interfaces.

### ***Rate Selectability 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” on page 345](#).

### ***Rate Selectability 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” on page 347](#).

### ***Rate Selectability on JNP10K-LC2101***

JNP10K-LC2101 is a fixed-configuration MPC and is supported on MX10008 routers. JNP10K-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 JNP10K-LC2101 MPCs. By default, each JNP10K-LC2101 MPC provides a maximum bandwidth of 1.44 Tbps. JNP10K-LC2101 has six Packet Forwarding Engines, each providing a maximum bandwidth of up to 240 Gbps, which cannot be oversubscribed. You can configure JNP10K-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 JNP10K-LC2101, the demand on each Packet Forwarding Engine must be less than or equal to 240 Gbps. However, if you have configured JNP10K-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 JNP10K-LC2101 MPC, see [“Interface Naming Conventions for JNP10K-LC2101” on page 352](#).

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* *gether-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).

### **Rate Selectability 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” on page 358](#).

### **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 Rate Selectability on MPC7E \(Multi-Rate\) to Enable Different Port Speeds | 389](#)

[Configuring Rate Selectability on MIC-MRATE to Enable Different Port Speeds | 385](#)

*Junos Continuity Software User Guide (Junos OS Release 14.1R4 and Later Releases)*

[Supported Active Physical Ports for Configuring Rate Selectability to Prevent Oversubscription | 364](#)

[Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds | 406](#)

## Guidelines for Configuring Rate Selectability

This topic describes the guidelines to consider when configuring rate selectability at the port level or the PIC or MIC level.

### ***Guidelines for Configuring Rate Selectability 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 Rate Selectability 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 Rate Selectability 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 Rate Selectability on MPC7E \(Multi-Rate\) to Enable Different Port Speeds](#) | 389

## MX10003 MPC Rate-Selectability Overview

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 Rate Selectability on MX10003 MPC to Enable Different Port Speeds” on page 394](#). 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 Rate Selectability on MX10003 MPC to Enable Different Port Speeds” on page 394](#).

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 31 on page 321](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 31: Rate Selectability 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 Rate Selectability to Prevent Oversubscription on MX10003 MPC” on page 367](#). 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* gigether-options speed 1g**

Table 32 on page 323 summarizes the port mode configuration at the Packet Forwarding Engine level.

**Table 32: 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 33 on page 323 summarizes the PIC mode configuration.

**Table 33: PIC Mode Configuration**

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



Table 33: PIC Mode Configuration (*continued*)

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

### 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 Rate Selectability on MX10003 MPC to Enable Different Port Speeds | 394](#)  
*speed (Gigabit Ethernet interface)*

### MX204 Router Rate-Selectability Overview

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 1-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” on page 369](#).

If your configuration does not comply to the following rules, it is an invalid configuration. In this case, the invalid configuration will not take in to effect and the router stays with the prior configuration:

- The MX204 routers contains one Packet Forwarding Engine that has a total of four 'MACs' (otherwise called as 'port groups'). Two MACs out of four are pre-configured as 1x100GE each and cannot be changed. The other two MACs can be channelized as one of the following each – either one 100-Gigabit

Ethernet ports, or three 40-Gigabit Ethernet ports (or one or more ports can be channelized down to 10-Gigabit Ethernet ports too), or 12 10-Gigabit Ethernet ports.

**NOTE:**

- One Gigabit Ethernet port can also be substituted for 10-Gigabit Ethernet ports.
- You cannot use the same MAC (or 'port groups') on both PICs.

- Two modes - PIC and port modes are supported and both PICs must be running in same mode:
  - PIC mode is the default mode that shares MAC across PICs, and all ports on both PICs run at the same speed. On MX204 routers contain 24 10-Gigabit Ethernet ports as all 4 ports on PIC 0 are channelized down to four 10-Gigabit Ethernet ports (16) and using all eight 10-Gigabit Ethernet ports on PIC 1.
  - Port mode allows you to run ports at different speeds.
  - You cannot spread or share MACs across PICs.

**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 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 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 34 on page 328](#) summarizes the rate selectability of the MX204 routers.

Table 34: 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 (QSFP-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 operation mode is supported with **speed 1g** configuration. For optics other than SFP-T, in 1G mode, the peer interfaces must be configured to non-autonegotiation mode.

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 Rate Selectability of MX204 Routers**

You can also configure rate selectability on MX204 routers.

[Table 35 on page 329](#) summarizes the user-configurable rate selectability of MX204 routers.

**Table 35: 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                                              |

Table 35: Configurable Rate Selectability of MX204 Router (*continued*)

| Port Speed Configuration on PIC 0 (Gbps) | Port Speed Configuration on PIC 1 (Gbps)            |
|------------------------------------------|-----------------------------------------------------|
| 40                                       | 0<br><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 36: 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><br>Which means 16 ports from PIC 0 and 8 Ports from PIC 1.      | 20<br><br>Which means 12 ports from PIC 0 and 8 Ports from PIC 1. |
| 40 Gigabit Ethernet Ports   | 4<br><br>Only 4 ports from PIC 0 as PIC 1 supports only 10 Gbps Speed. | 4                                                                 |
| 100 Gigabit Ethernet Ports  | 4<br><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 37: 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. |    |    |    |    |    |    |    |
| 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 |
| 100       | 10/40 | 10/40 | 10/40 | 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. |



| PIC Mode                                |    |    |    |                                                                    |
|-----------------------------------------|----|----|----|--------------------------------------------------------------------|
| PIC 0 (with four rate-selectable ports) |    |    |    | PIC 1 (with eight 10-Gigabit Ethernet ports)                       |
| 10                                      | 10 | 10 | 10 | 10<br><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* together-options speed 1g**

SEE ALSO

[Configuring Rate Selectability on MX204 to Enable Different Port Speeds | 398](#)

*speed*

*speed (Gigabit Ethernet interface)*

[show chassis pic | 1504](#)

*number-of-ports*

*pic-mode*

**MPC10E-15C-MRATE Rate-Selectability Overview**

MPC10E-15C-MRATE (15x100G) is a fixed-configuration Modular Port Concentrator (MPC) for MX240, MX480, and MX960 routers, that is capable of delivering up to 1.5T per-slot bandwidth. It consists of three PFE complexes per MPC and each PFE is capable of delivering up to 500G throughput. It supports three PICs per MPC, one PIC per PFE. Each PIC hosts five QSFP28 (5xQSFP28) ports that supports 4x10-Gbps, 40-Gbps, 100-Gbps, and 400-Gbps speeds using QSFP+, QSFP28, and QSFP56-DD optics types. The 4x10-Gbps speed is supported using breakout cables.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE.

For more information, see [“Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds” on page 406](#). 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 Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds” on page 406](#).

The ports on the MPC10E-15C-MRATE are called rate-selectable or multirate ports as they support multiple port speeds. You can choose to configure all ports in a PIC to operate at the same speed or configure all the ports at different supported speeds.

In MPC10E-15C-MRATE, you can choose to configure:

- A port in 4x10-Gigabit Ethernet mode using QSFP-4x10-Gigabit Ethernet optics and 4x10-Gigabit Ethernet breakout cables.
- A port in 40-Gigabit Ethernet mode using QSFP optics.
- A port in 100-Gigabit Ethernet mode using QSFP28 optics.
- A port (port 4) in 400-Gigabit Ethernet mode using QSFP56-DD optics.

**NOTE:**

- When you change the speed at PIC level, the PIC restarts automatically with the new configured speed.
- When you change the speed of a particular port explicitly (by port-level speed configuration), then only that particular port is reset automatically and other ports in that PIC remain unaffected.

The MPC10E-15C-MRATE supports three Packet Forwarding Engines (PFEs). The forwarding capacity of each Packet Forwarding Engine is 500Gbps which cannot be oversubscribed. The MPC10E-15C-MRATE supports 15 ports. Each PFE is mapped to 5 ports of the PIC. Use the command **show chassis pic fpc-slot slot-number pic-slot slot-number** to display PFE mapping information and port speed information.

[Table 38 on page 334](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

Table 38: Rate Selectability of MPC10E-15C-MRATE

| PIC              | Port Number | Port Speed Supported                                                                         | Remarks                                                                                                                                                                                                                                                                                                                                   |
|------------------|-------------|----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (or PFE 0) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet | <ul style="list-style-type: none"> <li>• 4x10GE support using QSFP-4x10GE breakout optics.</li> <li>• 1x40GE support using QSFP optics.</li> <li>• 1x100GE support using QSFP28 optics.</li> <li>• 1x400GE support using QSFP56-DD optics.</li> </ul> <p>NOTE: By default, all the active ports operate in 100-Gigabit Ethernet mode.</p> |
| PIC 1 (or PFE 1) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet |                                                                                                                                                                                                                                                                                                                                           |
| PIC 2 (or PFE 2) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet |                                                                                                                                                                                                                                                                                                                                           |

To view the speed configured for the interface, execute the *show interfaces extensive* command. The **Speed** output parameter in the command output indicates the current operation speed of the interface.

For example:

```

user@host>show interfaces et-4/1/0 extensive
Physical interface: et-4/1/0, Enabled, Physical link is Up
  Interface index: 220, SNMP ifIndex: 539, Generation: 223
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, 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
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags        : None
  CoS queues        : 8 supported, 8 maximum usable queues
  Schedulers        : 0
  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: 00:ca:fe:91:00:01, Hardware address: f4:a7:39:a8:b5:7e
Last flapped  : 2019-01-03 13:17:40 PST (00:02:43 ago)
Statistics last cleared: Never

...

```

In this example, the **Speed** output parameter displays **100Gbps**, which means the operation speed of et-4/1/0 interface is 100-Gbps.

Note the following important information while configuring rate selectability on the MPC10E-15C-MRATE:

- 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, PIC 1, and PIC 2) operate at 100-Gbps speed.
- When you configure rate selectability at the port level, only the configured ports are created in that PIC. Other ports are not created.
- You cannot configure rate selectability at the PIC level and the port level simultaneously. Commit fails when you try to apply such configuration.

#### ***User-Configurable Rate Selectability of MPC10E-15C-MRATE***

You can also configure rate selectability on MPC10E-15C-MRATE.

[Table 39 on page 335](#) summarizes the user-configurable rate selectability of MPC10E-15C-MRATE.

**Table 39: Port speed capability of MPC10E-15C-MRATE**

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

**NOTE:** Different PICs in the MPC10E-15C-MRATE can operate at different speeds. That is, PIC speed at one PIC has no bearing on the other PICs in the MPC.

SEE ALSO

[Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds](#) | 406

*speed (Gigabit Ethernet interface)*

## MPC10E-10C-MRATE Rate-Selectability Overview

MPC10E-10C-MRATE (10x100G) is a fixed-configuration Modular Port Concentrator (MPC) for MX240, MX480, and MX960 routers, that is capable of delivering up to 1Tbps per-slot bandwidth. It consists of two Packet Forwarding Engine (PFE) complexes per MPC and each PFE is capable of delivering up to 500G throughput. It supports two PICs per MPC, one PIC per PFE. Each PIC hosts five QSFP28 (5xQSFP28) ports that supports 4x10-Gbps, 40-Gbps, 100-Gbps, and 400-Gbps speeds using QSFP+, QSFP28, QSFP56-DD optics types. The 4x10-Gbps speed is supported using breakout cables.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE.

For more information, see [“Configuring Rate Selectability on MPC10E-10C-MRATE to Enable Different Port Speeds” on page 410](#). 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 Rate Selectability on MPC10E-10C-MRATE to Enable Different Port Speeds” on page 410](#).

The ports on the MPC10E-10C-MRATE are called rate-selectable or multirate ports as they support multiple port speeds. You can port choose to configure all ports in a PIC to operate at the same speed or configure all the ports at different supported speeds.

In MPC10E-10C-MRATE, you can choose to configure:

- A port in 4x10-Gigabit Ethernet mode using QSFP-4x10-Gigabit Ethernet optics and 4x10-Gigabit Ethernet breakout cables.
- A port in 40-Gigabit Ethernet mode using QSFP optics.
- A port in 100-Gigabit Ethernet mode using QSFP28 optics.
- A port (port 4) in 400-Gigabit Ethernet mode using QSFP56-DD optics.

### NOTE:

- When you change the speed at PIC level, the PIC restarts automatically with the new configured speed.
- When you change the speed of a particular port explicitly (by port-level speed configuration), then only that particular port is reset automatically and other ports in that PIC remain unaffected.

The MPC10E-10C-MRATE supports two Packet Forwarding Engines. The forwarding capacity of each Packet Forwarding Engine is 500Gbps which cannot be oversubscribed. The MPC10E-10C-MRATE supports 10 ports. Each PFE is mapped to 5 ports of the PIC. Use the command **show chassis pic fpc-slot slot-number pic-slot slot-number** to display PFE mapping information and port speed information.

[Table 40 on page 337](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

**Table 40: Rate Selectability of MPC10E-10C-MRATE**

| PIC              | Port Number | Port Speed Supported                                                                         | Remarks                                                                                                                                                                                            |
|------------------|-------------|----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (or PFE 0) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet | <ul style="list-style-type: none"> <li>• 4x10GE support using QSFP-4x10GE breakout optics.</li> <li>• 1x40GE support using QSFP optics.</li> <li>• 1x100GE support using QSFP28 optics.</li> </ul> |
| PIC 1 (or PFE 1) | 0-4         | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet<br>400-Gigabit Ethernet | <ul style="list-style-type: none"> <li>• 1x400GE support using QSFP56-DD optics.</li> </ul> <p><b>NOTE:</b> By default, all the active ports operate in 100-Gigabit Ethernet mode.</p>             |

To view the speed configured for the interface, execute the *show interfaces extensive* command. The **Speed** output parameter in the command output indicates the current operation speed of the interface.

For example:

```
user@host>show interfaces et-4/1/0 extensive
Physical interface: et-4/1/0, Enabled, Physical link is Up
  Interface index: 220, SNMP ifIndex: 539, Generation: 223
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, 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
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  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: 00:ca:fe:91:00:01, Hardware address: f4:a7:39:a8:b5:7e
Last flapped  : 2019-01-03 13:17:40 PST (00:02:43 ago)
Statistics last cleared: Never

...

```

In this example, the **Speed** output parameter displays **100Gbps**, which means the operation speed of et-4/1/0 interface is 100-Gbps.

Note the following important information while configuring rate selectability on the MPC10E-10C-MRATE:

- 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 configure rate selectability at the port level, only the configured ports are created in that PIC. Other ports are not created.
- You cannot configure rate selectability at the PIC level and the port level simultaneously. Commit fails when you try to apply such configuration.

#### ***User-Configurable Rate Selectability of MPC10E-10C-MRATE***

You can also configure rate selectability on MPC10E-10C-MRATE.

[Table 41 on page 338](#) summarizes the user-configurable rate selectability of MPC10E-10C-MRATE.

**Table 41: Port speed capability of MPC10E-10C-MRATE**

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

**NOTE:** Different PICs in the MPC10E-10C-MRATE can operate at different speeds. That is, PIC speed at one PIC has no bearing on the other PICs in the MPC.

## SEE ALSO

---

[Configuring Rate Selectability on MPC10E-10C-MRATE to Enable Different Port Speeds | 410](#)

*speed (Gigabit Ethernet interface)*

---

*MPC10E-15C-MRATE Overview*

---

[Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds | 406](#)



## PTX10003 Router Rate-Selectability 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:

| QSFP Transceiver | Native Port Speeds       | Channelization Options |
|------------------|--------------------------|------------------------|
| 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 | 1291](#)

[speed \(Ethernet\) | 1359](#)

[fpc](#)

[pic](#)

## MX2K-MPC11E Rate Selectability Overview

The MX2K-MPC11E is a fixed-configuration MPC that contains 8 built-in PICs, **PIC0** to **PIC7**. Each PIC has five physical ports that support quad small form-factor pluggable plus (QSFP28) transceivers. The default port speed is 100 Gbps for all ports. Only port 0 of **PIC0** to **PIC7** supports speeds of 4x10 Gbps (using breakout cables) and 40 Gbps, in addition to the default speed.

Each MX2K-MPC11E provides a maximum bandwidth of 4 Tbps. The MX2K-MPC11E has eight Packet Forwarding Engines, each providing a maximum bandwidth of up to 500 Gbps which cannot be oversubscribed. The ports on the MX2K-MPC11E are called rate-selectable or multi-rate ports as they support multiple port speeds. You can choose to configure all ports in a PIC to operate at the same speed or configure all the ports at different supported speeds.

On the MX2K-MPC11E, you can choose to configure:

- Port 0 of every PIC in 4x10-Gigabit Ethernet mode using QSFP-4x10-Gigabit Ethernet optics and 4x10-Gigabit Ethernet breakout cables.
- Port 0 of every PIC in 40-Gigabit Ethernet mode using QSFP optics.
- All other ports in every PIC in 100-Gigabit Ethernet mode using QSFP28 optics.

### NOTE:

- When you change the speed at PIC level, the PIC restarts automatically with the new configured speed.
- When you change the speed of a particular port in a given PIC explicitly (by port-level speed configuration), then only that particular port is affected and other ports in that PIC remain unaffected.

For information about how to configure rate selectability, see [“Configuring Rate Selectability on the MX2K-MPC11E to Enable Different Port Speeds” on page 414](#).

[Table 42 on page 343](#) summarizes the Packet Forwarding Engine mapping and the supported port speeds.

Table 42: Rate Selectability for the MX2K-MPC11E

| PIC              | Port Number | Port Speed Supported                                                 | Optics Supported                                                                                                                                                                                   |
|------------------|-------------|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PIC 0 (or PFE 0) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet | <ul style="list-style-type: none"> <li>• 4x10GE support using QSFP-4x10GE breakout optics.</li> <li>• 1x40GE support using QSFP optics.</li> <li>• 1x100GE support using QSFP28 optics.</li> </ul> |
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                                                                                                                                                                                                    |
| PIC 1 (or PFE 1) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet | <b>NOTE:</b> By default, all the active ports operate in 100-Gigabit Ethernet mode.                                                                                                                |
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                                                                                                                                                                                                    |
| PIC 2 (or PFE 2) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet |                                                                                                                                                                                                    |
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                                                                                                                                                                                                    |
| PIC 3 (or PFE 3) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet |                                                                                                                                                                                                    |
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                                                                                                                                                                                                    |
| PIC 4 (or PFE 4) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet |                                                                                                                                                                                                    |
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                                                                                                                                                                                                    |
| PIC 5 (or PFE 5) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet |                                                                                                                                                                                                    |

Table 42: Rate Selectability for the MX2K-MPC11E (continued)

| PIC              | Port Number | Port Speed Supported                                                 | Optics Supported |
|------------------|-------------|----------------------------------------------------------------------|------------------|
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                  |
| PIC 6 (or PFE 6) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet |                  |
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                  |
| PIC 7 (or PFE 7) | 0           | 40-Gigabit Ethernet<br>4x10-Gigabit Ethernet<br>100-Gigabit Ethernet |                  |
|                  | 1-4         | 100-Gigabit Ethernet                                                 |                  |

SEE ALSO

[Understanding Rate Selectability | 314](#)
[Understanding Interface Naming Conventions for MX2K-MPC11E MPC | 360](#)
**Release History Table**

| Release                | Description                                                                                                                                              |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 1-Gbps speed on 10-Gigabit Ethernet ports of the JNP10K-LC2101 MPC.                         |
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE.                                           |
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE.                                           |
| <a href="#">18.1R1</a> | 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. |
| <a href="#">18.1R1</a> | Starting with Junos OS Release 18.1R1, the 10-Gbps port can operate in 1-Gbps mode also.                                                                 |

## RELATED DOCUMENTATION

[Preventing Oversubscription Using Active Physical Ports | 364](#)

[Configuring Rate Selectability | 379](#)

## Interface Naming Conventions for Rate Selectability

### IN THIS SECTION

- [Interface Naming Conventions for MPC7E-MRATE | 345](#)
- [Interface Naming Conventions for MIC-MRATE | 347](#)
- [Interface Naming Conventions for MX10003 MPC | 350](#)
- [Interface Naming Conventions for JNP10K-LC2101 | 352](#)
- [Interface Naming Conventions for MPC10E-15C-MRATE MPC | 355](#)
- [Interface Naming Conventions for MPC10E-10C-MRATE MPC | 356](#)
- [Interface Naming Conventions for MIC-MACSEC-20GE | 358](#)
- [Understanding Interface Naming Conventions for MX2K-MPC11E MPC | 360](#)

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 rate selectability.

### 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 43 on page 346](#) lists the naming conventions for interfaces on MPC7E-MRATE for MX240, MX480, MX960, MX2010, and MX2020 routers.

**Table 43: 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>               | -                              |
|                          | <b>xe-x/1/4:[0-3]</b>         | <b>et-x/1/4</b>               | -                              |
|                          | <b>xe-x/1/5:[0-3]</b>         | <b>et-x/1/5</b>               | <b>et-x/1/5</b>                |

SEE ALSO

[Configuring Rate Selectability on MPC7E \(Multi-Rate\) to Enable Different Port Speeds](#) | 389

## 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 44 on page 348](#) lists the naming conventions used for interfaces on MIC-MRATE when installed on slot 0 of MPC8E and MPC9E. [Table 45 on page 349](#) 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 44: 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-x/0/0:[0-3]</b>         | <b>et-x/0/0</b>               | <b>et-x/0/0</b>                |
|                          | <b>xe-x/0/1:[0-3]</b>         | <b>et-x/0/1</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>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>               | -                              |
| 1                        | <b>xe-x/0/6:[0-3]</b>         | <b>et-x/0/6</b>               | <b>et-x/0/6</b>                |
|                          | <b>xe-x/0/7:[0-3]</b>         | <b>et-x/0/7</b>               | <b>et-x/0/7</b>                |
|                          | <b>xe-x/0/8:[0-3]</b>         | <b>et-x/0/8</b>               | <b>et-x/0/8</b>                |
|                          | <b>xe-x/0/9:[0-3]</b>         | <b>et-x/0/9</b>               | <b>et-x/0/9</b>                |
|                          | <b>xe-x/0/10:[0-3]</b>        | <b>et-x/0/10</b>              | -                              |
|                          | <b>xe-x/0/11:[0-3]</b>        | <b>et-x/0/11</b>              | -                              |

Table 45: 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                        | xe-x/1/0:[0-3]                | et-x/1/0                      | et-x/1/0                       |
|                          | xe-x/1/1:[0-3]                | et-x/1/1                      | et-x/1/1                       |
|                          | xe-x/1/2:[0-3]                | et-x/1/2                      | et-x/1/2                       |
|                          | xe-x/1/3:[0-3]                | et-x/1/3                      | et-x/1/3                       |
|                          | xe-x/1/4:[0-3]                | et-x/1/4                      | -                              |
|                          | xe-x/1/5:[0-3]                | et-x/1/5                      | -                              |
| 3                        | xe-x/1/6:[0-3]                | et-x/1/6                      | et-x/1/6                       |
|                          | xe-x/1/7:[0-3]                | et-x/1/7                      | et-x/1/7                       |
|                          | xe-x/1/8:[0-3]                | et-x/1/8                      | et-x/1/8                       |
|                          | xe-x/1/9:[0-3]                | et-x/1/9                      | et-x/1/9                       |
|                          | xe-x/1/10:[0-3]               | et-x/1/10                     | -                              |
|                          | xe-x/1/11:[0-3]               | et-x/1/11                     | -                              |

Table 46 on page 350 lists the naming conventions used for interfaces on MIC-MRATE when installed on slot 0 of MX10003 MPC.

Table 46: Interface Naming Convention for MIC-MRATE Installed on Slot 0 of Mx10003MPC

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 0                        | xe-x/0/0:[0-3]                | et-x/0/0                      | et-x/0/0                       |
|                          | xe-x/0/1:[0-3]                | et-x/0/1                      | et-x/0/1                       |
|                          | xe-x/0/2:[0-3]                | et-x/0/2                      | et-x/0/2                       |
|                          | xe-x/0/3:[0-3]                | et-x/0/3                      | et-x/0/3                       |
|                          | xe-x/0/4:[0-3]                | et-x/0/4                      | -                              |
|                          | xe-x/0/5:[0-3]                | et-x/0/5                      | -                              |
| 1                        | xe-x/0/6:[0-3]                | et-x/0/6                      | et-x/0/6                       |
|                          | xe-x/0/7:[0-3]                | et-x/0/7                      | et-x/0/7                       |
|                          | xe-x/0/8:[0-3]                | et-x/0/8                      | et-x/0/8                       |
|                          | xe-x/0/9:[0-3]                | et-x/0/9                      | et-x/0/9                       |
|                          | xe-x/0/10:[0-3]               | et-x/0/10                     | -                              |
|                          | xe-x/0/11:[0-3]               | et-x/0/11                     | -                              |

SEE ALSO

[Configuring Rate Selectability on MIC-MRATE to Enable Different Port Speeds | 385](#)

## 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 47 on page 351](#) lists the naming conventions used for interfaces on the fixed-port PIC when installed in slot 0 of the MX10003 MPC. [Table 48 on page 352](#) lists the naming conventions used for interfaces on the modular MIC when installed in slot 1 of the MPC.

**Table 47: Interface Naming Convention for the Fixed-Port PIC Installed in Slot 0 of MX10003 MPC**

| 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>               | –                              |
| 1                        | <b>xe-x/0/2:[0-3]</b>         | <b>et-x/0/2</b>               | –                              |
|                          | <b>xe-x/0/3:[0-3]</b>         | <b>et-x/0/3</b>               | –                              |
| 2                        | <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>               | –                              |

Table 48: Interface Naming Convention for Modular MIC Installed in Slot 1 of MX10003 MPC

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 0                        | xe-x/1/0:[0-3]                | et-x/1/0                      | et-x/1/0                       |
|                          | xe-x/1/1:[0-3]                | et-x/1/1                      | et-x/1/1                       |
|                          | xe-x/1/2:[0-3]                | et-x/1/2                      | et-x/1/2                       |
|                          | xe-x/1/3:[0-3]                | et-x/1/3                      | et-x/1/3                       |
| 1                        | xe-x/1/4:[0-3]                | et-x/1/4                      | et-x/1/4                       |
|                          | xe-x/1/5:[0-3]                | et-x/1/5                      | et-x/1/5                       |
|                          | xe-x/1/6:[0-3]                | et-x/1/6                      | et-x/1/6                       |
|                          | xe-x/1/7:[0-3]                | et-x/1/7                      | et-x/1/7                       |
| 2                        | xe-x/1/8:[0-3]                | et-x/1/8                      | et-x/1/8                       |
|                          | xe-x/1/9:[0-3]                | et-x/1/9                      | et-x/1/9                       |
|                          | xe-x/1/10:[0-3]               | et-x/1/10                     | et-x/1/10                      |
|                          | xe-x/1/11:[0-3]               | et-x/1/11                     | et-x/1/11                      |

SEE ALSO

*MX10003 MPC on MX10003 Router Overview*

## Interface Naming Conventions for JNP10K-LC2101

JNP10K-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.

JNP10K-LC2101 has a forwarding capacity of 240 Gbps for each Packet Forwarding Engine.

JNP10K-LC2101 has six Packet Forwarding Engines. In Junos OS Release 18.2R1 and later, you can upgrade JNP10K-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, JNP10K-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 JNP10K-LC2101 MPC follow the naming convention **et-fpc-slot/pic-slot/port-number**. The 10-Gigabit Ethernet interfaces configured on the JNP10K-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 JNP10K-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 JNP10K-LC2101 MPC that is installed in the MPC slot 0.

**NOTE:** Each Packet Forwarding Engine maps to a single built-in PIC on the JNP10K-LC2101.

Table 49 on page 353 lists the naming conventions used for interfaces on JNP10K-LC2101 for MX10008 routers. MX10008 routers support 8 JNP10K-LC2101 MPCs.

**NOTE:** The **x** in **et-x/0/0** and **xe-x/0/0:[0-3]** refers to the MPC slot number.

**Table 49: Interface Naming Convention for JNP10K-LC2101 MPC**

| 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>et-x/0/0</b>                |
|                          | <b>xe-x/0/1:[0-3]</b>         | <b>et-x/0/1</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>et-x/0/3</b>                |

Table 49: Interface Naming Convention for JNP10K-LC2101 MPC (continued)

| Packet Forwarding Engine | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|--------------------------|-------------------------------|-------------------------------|--------------------------------|
| 1                        | xe-x/1/0:[0-3]                | et-x/1/0                      | et-x/1/0                       |
|                          | xe-x/1/1:[0-3]                | et-x/1/1                      | et-x/1/1                       |
|                          | xe-x/1/2:[0-3]                | et-x/1/2                      | et-x/1/2                       |
|                          | xe-x/1/3:[0-3]                | et-x/1/3                      | et-x/1/3                       |
| 2                        | xe-x/2/0:[0-3]                | et-x/2/0                      | et-x/2/0                       |
|                          | xe-x/2/1:[0-3]                | et-x/2/1                      | et-x/2/1                       |
|                          | xe-x/2/2:[0-3]                | et-x/2/2                      | et-x/2/2                       |
|                          | xe-x/2/3:[0-3]                | et-x/2/3                      | et-x/2/3                       |
| 3                        | xe-x/3/0:[0-3]                | et-x/3/0                      | et-x/3/0                       |
|                          | xe-x/3/1:[0-3]                | et-x/3/1                      | et-x/3/1                       |
|                          | xe-x/3/2:[0-3]                | et-x/3/2                      | et-x/3/2                       |
|                          | xe-x/3/3:[0-3]                | et-x/3/3                      | et-x/3/3                       |
| 4                        | xe-x/4/0:[0-3]                | et-x/4/0                      | et-x/4/0                       |
|                          | xe-x/4/1:[0-3]                | et-x/4/1                      | et-x/4/1                       |
|                          | xe-x/4/2:[0-3]                | et-x/4/2                      | et-x/4/2                       |
|                          | xe-x/4/3:[0-3]                | et-x/4/3                      | et-x/4/3                       |
| 5                        | xe-x/5/0:[0-3]                | et-x/5/0                      | et-x/5/0                       |
|                          | xe-x/5/1:[0-3]                | et-x/5/1                      | et-x/5/1                       |
|                          | xe-x/5/2:[0-3]                | et-x/5/2                      | et-x/5/2                       |
|                          | xe-x/5/3:[0-3]                | et-x/5/3                      | et-x/5/3                       |

SEE ALSO

| *JNP10K-LC2101 MPC on MX10008 Routers Overview*

## Interface Naming Conventions for MPC10E-15C-MRATE MPC

Starting with Junos OS Release 19.1R1, the MPC10E-15C-MRATE (15x100G) is introduced for MX240, MX480, MX960 platforms, that is capable of delivering up to 1.5 terabits per-slot bandwidth. The MPC10E-15C-MRATE consists of three Packet Forwarding Engines (PFE) per MPC and each PFE is capable of delivering up to 500Gbps throughput. It supports three PICs per MPC, that is, one PIC per PFE, and five ports per PIC (that is, fifteen ports per MPC). Each PIC contains 5xQSFP28 ports that supports 400-Gbps, 100-Gbps, 40-Gbps, and 10-Gbps speeds using QSFP56-DD, QSFP28, and QSFP+ optics. The 10-Gbps speed is supported by using break out cables.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE.

### NOTE:

- MPC10E-15C-MRATE is only supported on MX240, MX480, MX960 routers and not on MX2000 series and MX10000 lines of routers.
- If you install the MPC10E-15C-MRATE, then the Switch Control Board MX-SCBE3 must be installed in the router in order for the MPC10E-15C-MRATE to come online.
- By default, the MPC10E-15C-MRATE ports are configured as 100-Gigabit Ethernet ports.

The interface names for the 10-Gigabit Ethernet interfaces uses the prefix xe- while the interface names for 25, 40 and 100 interfaces uses the prefix et-. When multiple interfaces (or IFDs) are supported in a physical port, the colon (:) notation is used in the interface naming conventions (Example: 4x10GE interfaces). The colon (:) is used as a delimiter to differentiate the multiple interfaces (or IFDs) in a physical port.

By default, the 400-Gigabit Ethernet, 100-Gigabit Ethernet, and 40-Gigabit Ethernet interfaces appear in the et-fpc/pic/port format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the xe-fpc/pic/port:channel format, where channel can be a value of 0 through 3.

**NOTE:** In the interface naming convention et-x/y/z and xe-x/y/z:[0-3] :

- x refers to the FPC slot number.
- y refers to the PIC slot number. The valid range is 0 to 2.
- z refers to the physical port number. The valid range is 0 to 4.



Table 50: Interface Naming Convention for MPC10E-15C-MRATE

| PIC | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface | 400-Gigabit Ethernet Interface |
|-----|-------------------------------|-------------------------------|--------------------------------|--------------------------------|
| 0   | xe-x/0/0:[0-3]                | et-x/0/0                      | et-x/0/0                       | -                              |
|     | xe-x/0/1:[0-3]                | et-x/0/1                      | et-x/0/1                       | -                              |
|     | xe-x/0/2:[0-3]                | et-x/0/2                      | et-x/0/2                       | -                              |
|     | xe-x/0/3:[0-3]                | et-x/0/3                      | et-x/0/3                       | -                              |
|     | xe-x/0/4:[0-3]                | et-x/0/4                      | et-x/0/4                       | et-x/0/4                       |
| 1   | xe-x/1/0:[0-3]                | et-x/1/0                      | et-x/1/0                       | -                              |
|     | xe-x/1/1:[0-3]                | et-x/1/1                      | et-x/1/1                       | -                              |
|     | xe-x/1/2:[0-3]                | et-x/1/2                      | et-x/1/2                       | -                              |
|     | xe-x/1/3:[0-3]                | et-x/1/3                      | et-x/1/3                       | -                              |
|     | xe-x/1/4:[0-3]                | et-x/1/4                      | et-x/1/4                       | et-x/0/4                       |
| 2   | xe-x/2/0:[0-3]                | et-x/2/0                      | et-x/2/0                       | -                              |
|     | xe-x/2/1:[0-3]                | et-x/2/1                      | et-x/2/1                       | -                              |
|     | xe-x/2/2:[0-3]                | et-x/2/2                      | et-x/2/2                       | -                              |
|     | xe-x/2/3:[0-3]                | et-x/2/3                      | et-x/2/3                       | -                              |
|     | xe-x/2/4:[0-3]                | et-x/2/4                      | et-x/2/4                       | et-x/0/4                       |

SEE ALSO

[Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds | 406](#)

### Interface Naming Conventions for MPC10E-10C-MRATE MPC

Starting with Junos OS Release 19.2R1, the MPC10E-10C-MRATE (10x100G) is introduced for MX240, MX480, MX960 platforms, that is capable of delivering up to 1 terabits per-slot bandwidth. The

MPC10E-10C-MRATE consists of two Packet Forwarding Engines (PFE) per MPC and each PFE is capable of delivering up to 500Gbps throughput. It supports two PICs per MPC, that is, one PIC per PFE, and five ports per PIC (that is, ten ports per MPC). Each PIC contains 5xQSFP28 ports that supports 400-Gbps, 100-Gbps, 40-Gbps, and 10-Gbps speeds using QSFP56-DD, QSFP28 and QSFP+ optics. The 10-Gbps speed is supported by using break out cables.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE.

**NOTE:**

- MPC10E-10C-MRATE is only supported on MX240, MX480, MX960 routers and not on MX2000 series and MX10000 lines of routers.
- If you install the MPC10E-10C-MRATE, then the Switch Control Board MX-SCBE3 must be installed in the router for the MPC10E-10C-MRATE to come online.
- By default, the MPC10E-10C-MRATE ports are configured as 100-Gigabit Ethernet ports.

The interface names for the 10-Gigabit Ethernet interfaces uses the prefix `xe-` while the interface names for 40, 100, and 400 interfaces uses the prefix `et-`. When multiple interfaces (or IFDs) are supported in a physical port, the colon (:) notation is used in the interface naming conventions (Example: 4x10GE interfaces). The colon (:) is used as a delimiter to differentiate the multiple interfaces (or IFDs) in a physical port.

By default, the 400-Gigabit Ethernet, 100-Gigabit Ethernet, and 40-Gigabit Ethernet interfaces appear in the `et-fpc/pic/port` format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the `xe-fpc/pic/port:channel` format, where channel can be a value of 0 through 3.

**NOTE:** In the interface naming convention `et-x/y/z` and `xe-x/y/z:[0-3]` :

- x refers to the FPC slot number.
- y refers to the PIC slot number. The valid range is 0 to 1.
- z refers to the physical port number. The valid range is 0 to 4.

Table 51: Interface Naming Convention for MPC10E-10C-MRATE

| PIC | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface | 400-Gigabit Ethernet Interface |
|-----|-------------------------------|-------------------------------|--------------------------------|--------------------------------|
| 0   | xe-x/0/0:[0-3]                | et-x/0/0                      | et-x/0/0                       | -                              |
|     | xe-x/0/1:[0-3]                | et-x/0/1                      | et-x/0/1                       | -                              |
|     | xe-x/0/2:[0-3]                | et-x/0/2                      | et-x/0/2                       | -                              |
|     | xe-x/0/3:[0-3]                | et-x/0/3                      | et-x/0/3                       | -                              |
|     | xe-x/0/4:[0-3]                | et-x/0/4                      | et-x/0/4                       | et-x/0/4                       |
| 1   | xe-x/1/0:[0-3]                | et-x/1/0                      | et-x/1/0                       | -                              |
|     | xe-x/1/1:[0-3]                | et-x/1/1                      | et-x/1/1                       | -                              |
|     | xe-x/1/2:[0-3]                | et-x/1/2                      | et-x/1/2                       | -                              |
|     | xe-x/1/3:[0-3]                | et-x/1/3                      | et-x/1/3                       | -                              |
|     | xe-x/1/4:[0-3]                | et-x/1/4                      | et-x/1/4                       | et-x/1/4                       |

## SEE ALSO

[Configuring Rate Selectability on MPC10E-10C-MRATE to Enable Different Port Speeds | 410](#)

[MPC10E-15C-MRATE Overview](#)

[Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds | 406](#)

## 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 52: Interface Naming Convention for MIC-MACSEC-20GE**

| PIC   | 1-Gigabit Ethernet Interface | 10-Gigabit Ethernet Interface |
|-------|------------------------------|-------------------------------|
| PIC 0 | <code>ge-x/0/[0-9]</code>    | <code>xe-x/0/0</code>         |
| PIC 1 | <code>ge-x/1/[0-9]</code>    | <code>xe-x/1/0</code>         |
| PIC 2 | <code>ge-x/2/[0-9]</code>    | <code>xe-x/2/0</code>         |
| PIC 3 | <code>ge-x/3/[0-9]</code>    | <code>xe-x/3/0</code>         |

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 MX Series Routers*

*cipher-suite*

*MPC1 on MX Series Routers Overview*

*MPC2 on MX Series Routers Overview*

*MPC3E on MX Series Routers Overview*

**Understanding Interface Naming Conventions for MX2K-MPC11E MPC**

Starting with Junos OS Release 19.3R2, the MX2K-MPC11E is introduced for MX2010 and MX2020 routers. The MX2K-MPC11E is a fixed-configuration model port concentrator (MPC) which delivers bandwidth up to 4 Tbps per MPC slot. The MX2K-MPC11E has eight packet forwarding engines, each capable of delivering up to 500 Gbps throughput. The MX2K-MPC11E supports eight PICs—one PIC per PFE. There are five ports per PIC for a total of 40 ports per MX2K-MPC11E. Each PIC contains 5xQSFP28 ports. The default speed for all ports is 100-Gbps. Port 0 of each PIC supports 100-Gbps, 40-Gbps, and 4x10-Gbps speeds using QSFP28, QSFP+ optics. The 4x10-Gbps speed is supported by using break out cables.

**NOTE:**

- The MX2K-MPC11E supports only MX2010 and MX2020 routers.
- The MX2K-MPC11E supports only MX2000-SFB3 switch fabric boards. It does not support SFB1 or SFB2. The router must have at least one MX2000-SFB3 to power on.
- By default, the MX2K-MPC11E ports are configured as 100-Gigabit Ethernet ports.

The interface names for the 10-Gigabit Ethernet interfaces uses the prefix xe- while the interface names for 40-Gigabit Ethernet interfaces and 100-Gigabit Ethernet interfaces uses the prefix et-. When multiple interfaces are supported on a physical port, the colon (:) notation is used in the interface naming conventions

(Example: 4x10GE interfaces). The colon (:) is used as a delimiter to differentiate the multiple interfaces on a physical port.

By default, the 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces appear in the et-fpc/pic/port format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the xe-fpc/pic/port:channel format, where channel can be a value of 0 through 3.

**NOTE:** In the interface naming convention et-x/y/z and xe-x/y/z:[0-3] :

- x refers to the FPC slot number. The valid range is 0 through 19.
- y refers to the PIC slot number. The valid range is 0 through 7.
- z refers to the physical port number. The valid range is 0 through 4.

Table 53 on page 361 lists the interface naming conventions for the MX2K-MPC11E.

**Table 53: Interface Naming Convention for MPC11E**

| PIC | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|-----|-------------------------------|-------------------------------|--------------------------------|
| 0   | xe-x/0/0:[0-3]                | et-x/0/0                      | et-x/0/0                       |
|     |                               |                               | et-x/0/1                       |
|     |                               |                               | et-x/0/2                       |
|     |                               |                               | et-x/0/3                       |
|     |                               |                               | et-x/0/4                       |
| 1   | xe-x/1/0:[0-3]                | et-x/1/0                      | et-x/1/0                       |
|     |                               |                               | et-x/1/1                       |
|     |                               |                               | et-x/1/2                       |
|     |                               |                               | et-x/1/3                       |
|     |                               |                               | et-x/1/4                       |

Table 53: Interface Naming Convention for MPC11E (continued)

| PIC | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|-----|-------------------------------|-------------------------------|--------------------------------|
| 2   | xe-x/2/0:[0-3]                | et-x/2/0                      | et-x/2/0                       |
|     |                               |                               | et-x/2/1                       |
|     |                               |                               | et-x/2/2                       |
|     |                               |                               | et-x/2/3                       |
|     |                               |                               | et-x/2/4                       |
| 3   | xe-x/3/0:[0-3]                | et-x/3/0                      | et-x/3/0                       |
|     |                               |                               | et-x/3/1                       |
|     |                               |                               | et-x/3/2                       |
|     |                               |                               | et-x/3/3                       |
|     |                               |                               | et-x/3/4                       |
| 4   | xe-x/4/0:[0-3]                | et-x/4/0                      | et-x/4/0                       |
|     |                               |                               | et-x/4/1                       |
|     |                               |                               | et-x/4/2                       |
|     |                               |                               | et-x/4/3                       |
|     |                               |                               | et-x/4/4                       |
| 5   | xe-x/5/0:[0-3]                | et-x/5/0                      | et-x/5/0                       |
|     |                               |                               | et-x/5/1                       |
|     |                               |                               | et-x/5/2                       |
|     |                               |                               | et-x/5/3                       |
|     |                               |                               | et-x/5/4                       |

Table 53: Interface Naming Convention for MPC11E (continued)

| PIC | 10-Gigabit Ethernet Interface | 40-Gigabit Ethernet Interface | 100-Gigabit Ethernet Interface |
|-----|-------------------------------|-------------------------------|--------------------------------|
| 6   | xe-x/6/0:[0-3]                | et-x/6/0                      | et-x/6/0                       |
|     |                               |                               | et-x/6/1                       |
|     |                               |                               | et-x/6/2                       |
|     |                               |                               | et-x/6/3                       |
|     |                               |                               | et-x/6/4                       |
| 7   | xe-x/7/0:[0-3]                | et-x/7/0                      | et-x/7/0                       |
|     |                               |                               | et-x/7/1                       |
|     |                               |                               | et-x/7/2                       |
|     |                               |                               | et-x/7/3                       |
|     |                               |                               | et-x/7/4                       |

SEE ALSO

[Configuring Rate Selectability on the MX2K-MPC11E to Enable Different Port Speeds | 414](#)

**Release History Table**

| Release                | Description                                                                                                    |
|------------------------|----------------------------------------------------------------------------------------------------------------|
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE. |
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE. |

**RELATED DOCUMENTATION**

[Configuring Rate Selectability | 379](#)



## Preventing Oversubscription Using Active Physical Ports

### IN THIS SECTION

- [Supported Active Physical Ports for Configuring Rate Selectability to Prevent Oversubscription | 364](#)
- [Supported Active Physical Ports for Configuring Rate Selectability to Prevent Oversubscription on MX10003 MPC | 367](#)
- [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router | 369](#)
- [Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on MPC10E-15C-MRATE | 373](#)
- [Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on MPC10E-10C-MRATE | 376](#)

When oversubscription of Packet Forwarding Capacity is not supported, the demand on each PFE should be less than or equal to its forwarding capacity. To prevent oversubscription, you can configure the number of active ports that operate at the configured speed. Interfaces are created only for active ports. Use this topic for information about the supported active ports on specific line cards.

### Supported Active Physical Ports for Configuring Rate Selectability to Prevent Oversubscription

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, on MPC7E-MRATE, the demand on each Packet Forwarding Engine must be less than or equal to 240 Gbps.

**NOTE:** By default, the MPC7E-MRATE ports are configured as 10-Gigabit Ethernet ports.

When you configure rate selectability at the MIC level, all the ports supporting that port speed are enabled by default. This can lead to fabric oversubscription in certain cases. To prevent fabric oversubscription, you can configure the number of active ports that operate at the configured speed by using the **number-of-ports number-of-active-physical-ports** configuration statement. Additionally, interfaces are created only for the active ports.

**NOTE:** You cannot configure the number of active ports when you configure rate selectability at the port level.

Table 54 on page 365 lists the active physical ports on MPC7E-MRATE.

**Table 54: Active Physical Ports on MPC7E-MRATE MPC for Configuring Rate Selectability at PIC Level**

| Ports Configured<br>(number-of-ports<br>Statement) | Active Physical Ports for Different Configured Speeds |                     |                      |
|----------------------------------------------------|-------------------------------------------------------|---------------------|----------------------|
|                                                    | 10-Gigabit Ethernet                                   | 40-Gigabit Ethernet | 100-Gigabit Ethernet |
| 1                                                  | 0                                                     | 0                   | 2                    |
| 2                                                  | 0, 1                                                  | 0, 1                | 2, 5                 |
| 3                                                  | 0, 1, 2                                               | 0, 1, 2             | 2, 5                 |
| 4                                                  | 0, 1, 2, 3                                            | 0, 1, 2, 3          | 2, 5                 |
| 5                                                  | 0, 1, 2, 3, 4                                         | 0, 1, 2, 3, 4       | 2, 5                 |
| 6                                                  | 0, 1, 2, 3, 4, 5                                      | 0, 1, 2, 3, 4, 5    | 2, 5                 |

MPC8E has a forwarding capacity of 240 Gbps for each Packet Forwarding Engine. MPC9E has a forwarding capacity of 400 Gbps for 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, on MPC8E, the demand on each Packet Forwarding Engine must be less than or equal to 240 Gbps and on MPC9E, the demand on each Packet Forwarding Engine must be less than or equal to 400 Gbps.

**NOTE:** By default, the MIC-MRATE ports are configured as 10-Gigabit Ethernet ports.

Table 55 on page 366, Table 56 on page 366 list the active physical ports on MPC8E and MPC9E.

**Table 55: Active Physical Ports on MIC-MRATE on MPC8E MPC for Configuring Rate Selectability at MIC Level**

| Ports Configured<br>(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, 6                                                  | 0, 6                                 | 0, 6                 |
| 3                                               | 0, 1, 6                                               | 0, 1, 6                              | 0, 1, 6              |
| 4                                               | 0, 1, 6, 7                                            | 0, 1, 6, 7                           | 0, 1, 6, 7           |
| 5                                               | 0, 1, 2, 6, 7                                         | 0, 1, 2, 6, 7                        | 0, 1, 6, 7           |
| 6                                               | 0, 1, 2, 6, 7, 8                                      | 0, 1, 2, 6, 7, 8                     | 0, 1, 6, 7           |
| 7                                               | 0, 1, 2, 3, 6, 7, 8                                   | 0, 1, 2, 3, 6, 7, 8                  | 0, 1, 6, 7           |
| 8                                               | 0, 1, 2, 3, 6, 7, 8, 9                                | 0, 1, 2, 3, 6, 7, 8, 9               | 0, 1, 6, 7           |
| 9                                               | 0, 1, 2, 3, 4, 6, 7, 8, 9                             | 0, 1, 2, 3, 4, 6, 7, 8, 9            | 0, 1, 6, 7           |
| 10                                              | 0, 1, 2, 3, 4, 6, 7, 8, 9, 10                         | 0, 1, 2, 3, 4, 6, 7, 8, 9, 10        | 0, 1, 6, 7           |
| 11                                              | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10                      | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10     | 0, 1, 6, 7           |
| 12                                              | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11                  | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11 | 0, 1, 6, 7           |

**Table 56: Active Physical Ports on MIC-MRATE on MPC9E MPC and MPC8E MPC in 1.6T Mode for Configuring Rate Selectability at MIC Level**

| Ports Configured<br>(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, 6                                                  | 0, 6                | 0, 6                 |
| 3                                               | 0, 1, 6                                               | 0, 1, 6             | 0, 1, 6              |

**Table 56: Active Physical Ports on MIC-MRATE on MPC9E MPC and MPC8E MPC in 1.6T Mode for Configuring Rate Selectability at MIC Level (continued)**

| Ports Configured<br>( <b>number-of-ports</b> Statement) | Active Physical Ports for Different Configured Speeds |                                      |                        |
|---------------------------------------------------------|-------------------------------------------------------|--------------------------------------|------------------------|
|                                                         | 10-Gigabit Ethernet                                   | 40-Gigabit Ethernet                  | 100-Gigabit Ethernet   |
| 4                                                       | 0, 1, 6, 7                                            | 0, 1, 6, 7                           | 0, 1, 6, 7             |
| 5                                                       | 0, 1, 2, 6, 7                                         | 0, 1, 2, 6, 7                        | 0, 1, 2, 6, 7          |
| 6                                                       | 0, 1, 2, 6, 7, 8                                      | 0, 1, 2, 6, 7, 8                     | 0, 1, 2, 6, 7, 8       |
| 7                                                       | 0, 1, 2, 3, 6, 7, 8                                   | 0, 1, 2, 3, 6, 7, 8                  | 0, 1, 2, 3, 6, 7, 8    |
| 8                                                       | 0, 1, 2, 3, 6, 7, 8, 9                                | 0, 1, 2, 3, 6, 7, 8, 9               | 0, 1, 2, 3, 6, 7, 8, 9 |
| 9                                                       | 0, 1, 2, 3, 4, 6, 7, 8, 9                             | 0, 1, 2, 3, 4, 6, 7, 8, 9            | 0, 1, 2, 3, 6, 7, 8, 9 |
| 10                                                      | 0, 1, 2, 3, 4, 6, 7, 8, 9, 10                         | 0, 1, 2, 3, 4, 6, 7, 8, 9, 10        | 0, 1, 2, 3, 6, 7, 8, 9 |
| 11                                                      | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10                      | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10     | 0, 1, 2, 3, 6, 7, 8, 9 |
| 12                                                      | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11                  | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11 | 0, 1, 2, 3, 6, 7, 8, 9 |

SEE ALSO

[Understanding Rate Selectability | 314](#)

### **Supported Active Physical Ports for Configuring Rate Selectability to Prevent Oversubscription on MX10003 MPC**

When you configure rate selectability at the MIC level, all the ports supporting that port speed are enabled by default. This can lead to fabric oversubscription in certain cases. To prevent fabric oversubscription, you can configure the number of active ports that operate at the configured speed by using the **number-of-ports number-of-active-physical-ports** configuration statement. Additionally, interfaces are created only for the active ports.

**NOTE:** You cannot configure the number of active ports when you configure rate selectability at the port level.

Starting in Junos OS Release 17.3R1, the MX10003 MPC supports rate selectability to prevent oversubscription of the Packet Forwarding Engine bandwidth.

Table 57 on page 368 lists the active physical ports on MX10003 MPC for Configuring Rate Selectability at the MIC Level.

**Table 57: Active Physical Ports on the MX10003 MPC for configuring rate selectability at the MIC level**

| Ports Configured<br>(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 58 on page 369 list the active physical ports on MX10003 MPC for Configuring Rate Selectability at PIC Level.

**Table 58: Active Physical Ports on MX10003 MPC for configuring rate selectability at the PIC level**

| Ports Configured<br>(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    |

SEE ALSO

[Understanding Rate Selectability | 314](#)

*MX10003 MPC on MX10003 Router Overview*

### Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on MX204 Router

The maximum capacity of an MX204 router is 400 Gbps, which cannot be oversubscribed. In MX204, the network ports are available in two groups (referred to as PICs), with restrictions around the number and type of ports that can be configured without oversubscription.

Starting in Junos OS Release 17.4R1, the MX204 supports rate selectability to prevent oversubscription of the Packet Forwarding Engine bandwidth. The MX204 Packet Forwarding Engine has four 100-Gigabit Ethernet QSFP28 ports (referred to as PIC 0 ports) and eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports). Each of the PIC 0 ports can be used as either a 100-Gigabit Ethernet QSFP28 port or a 40-Gigabit Ethernet QSFP28 port, or they can be configured as four 10-Gigabit Ethernet ports (using a breakout cable).

If you configure rate selectability at the PIC level, all the ports supporting that port speed are enabled by default. This can lead to oversubscription in certain cases. To prevent the oversubscription, you can configure the number of active ports that operate at the configured speed by using the **number-of-ports**

***number-of-active-physical-ports*** configuration statement. Additionally, interfaces are created only for the active ports.

**NOTE:**

- You cannot configure the number of active ports when you configure rate selectability at the port level.

### ***Invalid Port Configuration***

You must try to avoid configuring ports that can lead to oversubscription.

Following is an example of an invalid configuration:

```
4x100GE + 8x10GE
```

If you try to commit an invalid configuration, the configuration gets committed, but the port is not activated. This is because Junos OS allows you to configure a port before a line card is inserted. You will get an error message in the output of the *show chassis alarms* command and also in the log messages.

**NOTE:** When you are in port configuration mode, all the ports are configured as 10-Gigabit Ethernet.

**NOTE:** When you configure the QSFP28 ports with multiple port speeds, it can lead to oversubscription. To fix the issue, you must disable the ports on **PIC 1** by using the **set chassis fpc 0 pic 1 number-of-ports 0** command so that **PIC 0** can utilize the full capacity of the Packet Forwarding Engine.

### ***Configuring Active Ports on MX204 Router with Rate Selectability***

[Table 59 on page 371](#) summarizes the active ports with *number-of-ports* configured but without any rate selectability configuration for an MX204 router. Because there is no rate selectability configured, the default speed is used in these cases.

Table 59: Active Physical Ports on the MX204 Router for Configuring Rate Selectability at PIC level

| PIC   | Number of Ports<br>( <del>number-of-ports</del><br>Statement) | Active Ports                          |                                       |                                        |
|-------|---------------------------------------------------------------|---------------------------------------|---------------------------------------|----------------------------------------|
|       |                                                               | PIC Level 10-Gigabit Ethernet Profile | PIC Level 40-Gigabit Ethernet Profile | PIC Level 100-Gigabit Ethernet Profile |
| PIC 0 | 0                                                             | -                                     | -                                     | -                                      |
|       | 1                                                             | 0                                     | -                                     | -                                      |
|       | 2                                                             | 0, 1                                  | -                                     | -                                      |
|       | 3                                                             | 0, 1, 2                               | -                                     | -                                      |
|       | 4                                                             | 0, 1, 2, 3                            | -                                     | -                                      |
| PIC 1 | 0                                                             | -                                     | -                                     | -                                      |
|       | 1                                                             | 0                                     | -                                     | -                                      |
|       | 2                                                             | 0, 1                                  | -                                     | -                                      |
|       | 3                                                             | 0, 1, 2                               | -                                     | -                                      |
|       | 4                                                             | 0, 1, 2, 3                            | -                                     | -                                      |
|       | 5                                                             | 0, 1, 2, 3, 4                         | -                                     | -                                      |
|       | 6                                                             | 0, 1, 2, 3, 4, 5                      | -                                     | -                                      |
|       | 7                                                             | 0, 1, 2, 3, 4, 5, 6                   | -                                     | -                                      |
|       | 8                                                             | 0, 1, 2, 3, 4, 5, 6, 7                | -                                     | -                                      |

Table 60 on page 372 summarizes the active ports without *number-of-ports* configured but with rate selectability at PIC-level configuration for an MX204 router.



Table 60: Without number-of-ports But with Rate Selectability at PIC Level for MX204 Router

| PIC   | Active Ports                  |                               |                                |
|-------|-------------------------------|-------------------------------|--------------------------------|
|       | PIC-Level 10-Gigabit Ethernet | PIC-Level 40-Gigabit Ethernet | PIC-Level 100-Gigabit Ethernet |
| PIC 0 | 0-3                           | 0-3                           | 0-3                            |
| PIC 1 | 0-7                           | -                             | -                              |

Table 61 on page 372 summarizes the active ports with *number-of-ports* configured and rate selectability at PIC-level configuration for an MX204 router.

Table 61: With number-of-ports Rate Selectability at PIC level for MX204 Router

| PIC   | Number of Ports (number-of-ports Statement) | Active Ports                  |                               |                                |
|-------|---------------------------------------------|-------------------------------|-------------------------------|--------------------------------|
|       |                                             | PIC-Level 10-Gigabit Ethernet | PIC-Level 40-Gigabit Ethernet | PIC-Level 100-Gigabit Ethernet |
| PIC 0 | 0                                           | -                             | -                             | -                              |
|       | 1                                           | 0                             | 0                             | 0                              |
|       | 2                                           | 0, 1                          | 0, 1                          | 0, 1                           |
|       | 3                                           | 0, 1, 2                       | 0, 1, 2                       | 0, 1, 2                        |
|       | 4                                           | 0, 1, 2, 3                    | 0, 1, 2, 3                    | 0, 1, 2, 3                     |

Table 61: With number-of-ports Rate Selectability at PIC level for MX204 Router (*continued*)

| PIC   | Number of Ports<br>( <i>number-of-ports</i> Statement) | Active Ports                  |                               |                                |
|-------|--------------------------------------------------------|-------------------------------|-------------------------------|--------------------------------|
|       |                                                        | PIC-Level 10-Gigabit Ethernet | PIC-Level 40-Gigabit Ethernet | PIC-Level 100-Gigabit Ethernet |
| PIC 1 | 0                                                      | -                             | -                             | -                              |
|       | 1                                                      | 0                             | -                             | -                              |
|       | 2                                                      | 0, 1                          | -                             | -                              |
|       | 3                                                      | 0, 1, 2                       | -                             | -                              |
|       | 4                                                      | 0, 1, 2, 3                    | -                             | -                              |
|       | 5                                                      | 0, 1, 2, 3, 4                 | -                             | -                              |
|       | 6                                                      | 0, 1, 2, 3, 4, 5              | -                             | -                              |
|       | 7                                                      | 0, 1, 2, 3, 4, 5, 6           | -                             | -                              |
|       | 8                                                      | 0, 1, 2, 3, 4, 5, 6, 7        | -                             | -                              |

## SEE ALSO

[MX204 Router Overview](#)
[MX204 Router Rate-Selectability Overview | 325](#)

### Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on MPC10E-15C-MRATE

When you configure speed at the PIC level, all the ports supporting that port speed are enabled by default. This could lead to Packet Forwarding Engine (PFE) oversubscription. To prevent oversubscription, configure the number of active ports that operate at the configured speed by using the **number-of-ports** *number-of-active-physical-ports* configuration statement. Additionally, interfaces are created only for the active ports.

**NOTE:** You cannot configure the number of active ports when you configure rate at the port level.

Starting in Junos OS Release 19.1R1, the MPC10E-15C-MRATE supports rate selectability to prevent oversubscription of the PFE bandwidth.

Table 62 on page 374 summarizes the active ports with **number-of-ports** and without any rate selectability configuration for MPC10E-15C-MRATE. Because there is no rate selectability configured, the default mode and no **pic-mode** configuration is applied.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE.

**Table 62: Active Ports with number-of-ports and without rate-selectability**

| Number of Active Ports | Active Ports with PIC level rate selectability |                     |                      |                      |
|------------------------|------------------------------------------------|---------------------|----------------------|----------------------|
|                        | 10-Gigabit Ethernet                            | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 1                      | -                                              | -                   | 0                    | -                    |
| 2                      | -                                              | -                   | 0, 1                 | -                    |
| 3                      | -                                              | -                   | 0, 1, 2              | -                    |
| 4                      | -                                              | -                   | 0, 1, 2, 3           | -                    |
| 5                      | -                                              | -                   | 0, 1, 2, 3, 4        | -                    |

Table 63 on page 374 list the active ports without **number-of-ports** and with rate selectability at PIC level for MPC10E-15C-MRATE.

**Table 63: Active Ports without number-of-ports and with speed configured at PIC level**

| PIC Type           | Active Ports with speed configured at PIC level |                     |                      |                      |
|--------------------|-------------------------------------------------|---------------------|----------------------|----------------------|
|                    | 10-Gigabit Ethernet                             | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 5-CP28 PIC (PIC 0) | 0, 1, 2, 3, 4                                   | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |

Table 63: Active Ports without number-of-ports and with speed configured at PIC level (continued)

| PIC Type           | Active Ports with speed configured at PIC level |                     |                      |                      |
|--------------------|-------------------------------------------------|---------------------|----------------------|----------------------|
|                    | 10-Gigabit Ethernet                             | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 5QFP28 PIC (PIC 1) | 0, 1, 2, 3, 4                                   | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |
| 5QFP28 PIC (PIC 2) | 0, 1, 2, 3, 4                                   | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |

Table 64 on page 375 list the active ports with **number-of-ports** and rate selectability at PIC level for MPC10E-15C-MRATE.

Table 64: Active Ports with number-of-ports and rate selectability at PIC level

| PIC Type           | Number of Active Ports | Active Ports with speed configured at PIC level |                     |                      |                      |
|--------------------|------------------------|-------------------------------------------------|---------------------|----------------------|----------------------|
|                    |                        | 10-Gigabit Ethernet                             | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 5QFP28 PIC (PIC 0) | 1                      | 0                                               | 0                   | 0                    | 4                    |
|                    | 2                      | 0, 1                                            | 0, 1                | 0, 1                 | 4                    |
|                    | 3                      | 0, 1, 2                                         | 0, 1, 2             | 0, 1, 2              | 4                    |
|                    | 4                      | 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        | 4                    |
| 5QFP28 PIC (PIC 1) | 1                      | 0                                               | 0                   | 0                    | 4                    |
|                    | 2                      | 0, 1                                            | 0, 1                | 0, 1                 | 4                    |
|                    | 3                      | 0, 1, 2                                         | 0, 1, 2             | 0, 1, 2              | 4                    |
|                    | 4                      | 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        | 4                    |

Table 64: Active Ports with number-of-ports and rate selectability at PIC level (*continued*)

| PIC Type              | Number of Active Ports | Active Ports with speed configured at PIC level |                     |                      |                      |
|-----------------------|------------------------|-------------------------------------------------|---------------------|----------------------|----------------------|
|                       |                        | 10-Gigabit Ethernet                             | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 500P2B<br>PIC (PIC 2) | 1                      | 0                                               | 0                   | 0                    | 4                    |
|                       | 2                      | 0, 1                                            | 0, 1                | 0, 1                 | 4                    |
|                       | 3                      | 0, 1, 2                                         | 0, 1, 2             | 0, 1, 2              | 4                    |
|                       | 4                      | 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        | 4                    |

## SEE ALSO

[Understanding Rate Selectability | 314](#)
[number-of-ports](#)
[number-of-sub-ports | 1291](#)
[pic-mode](#)
[speed](#)

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

When you configure speed at the PIC level, all the ports supporting that port speed are enabled by default. This could lead to Packet Forwarding Engine (PFE) oversubscription. To prevent oversubscription, configure the number of active ports that operate at the configured speed by using the **number-of-ports** **number-of-active-physical-ports** configuration statement. Additionally, interfaces are created only for the active ports.

**NOTE:** You cannot configure the number of active ports when you configure rate at the port level.

Starting in Junos OS Release 19.2R1, the MPC10E-10C-MRATE supports rate selectability to prevent oversubscription of the PFE bandwidth.

[Table 65 on page 377](#) summarizes the active ports with **number-of-ports** and without any rate selectability configuration for MPC10E-10C-MRATE. Because there is no rate selectability configured, the default mode and no **pic-mode** configuration is applied.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE.

**Table 65: Active Ports with number-of-ports and without rate-selectability**

| Number of Active Ports | Active Ports with PIC level rate selectability |                     |                      |                      |
|------------------------|------------------------------------------------|---------------------|----------------------|----------------------|
|                        | 10-Gigabit Ethernet                            | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 1                      | -                                              | -                   | 0                    | -                    |
| 2                      | -                                              | -                   | 0, 1                 | -                    |
| 3                      | -                                              | -                   | 0, 1, 2              | -                    |
| 4                      | -                                              | -                   | 0, 1, 2, 3           | -                    |
| 5                      | -                                              | -                   | 0, 1, 2, 3, 4        | -                    |

[Table 66 on page 377](#) lists the active ports without **number-of-ports** and with rate selectability at PIC level for MPC10E-10C-MRATE.

**Table 66: Active Ports without number-of-ports and with speed configured at PIC level**

| PIC Type             | Active Ports with speed configured at PIC level |                     |                      |                      |
|----------------------|-------------------------------------------------|---------------------|----------------------|----------------------|
|                      | 10-Gigabit Ethernet                             | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 5xQSFP28 PIC (PIC 0) | 0, 1, 2, 3, 4                                   | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |
| 5xQSFP28 PIC (PIC 1) | 0, 1, 2, 3, 4                                   | 0, 1, 2, 3, 4       | 0, 1, 2, 3, 4        | 4                    |

[Table 67 on page 378](#) lists the active ports with **number-of-ports** and rate selectability at PIC level for MPC10E-10C-MRATE.

Table 67: Active Ports with number-of-ports and rate selectability at PIC level

| PIC Type            | Number of Active Ports | Active Ports with speed configured at PIC level |                     |                      |                      |
|---------------------|------------------------|-------------------------------------------------|---------------------|----------------------|----------------------|
|                     |                        | 10-Gigabit Ethernet                             | 40-Gigabit Ethernet | 100-Gigabit Ethernet | 400-Gigabit Ethernet |
| 50GFP28 PIC (PIC 0) | 1                      | 0                                               | 0                   | 0                    | 4                    |
|                     | 2                      | 0, 1                                            | 0, 1                | 0, 1                 | 4                    |
|                     | 3                      | 0, 1, 2                                         | 0, 1, 2             | 0, 1, 2              | 4                    |
|                     | 4                      | 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        | 4                    |
| 50GFP28 PIC (PIC 1) | 1                      | 0                                               | 0                   | 0                    | 4                    |
|                     | 2                      | 0, 1                                            | 0, 1                | 0, 1                 | 4                    |
|                     | 3                      | 0, 1, 2                                         | 0, 1, 2             | 0, 1, 2              | 4                    |
|                     | 4                      | 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        | 4                    |

## SEE ALSO

[Understanding Rate Selectability | 314](#)
[number-of-ports](#)
[number-of-sub-ports | 1291](#)
[pic-mode](#)
[speed](#)
[MPC10E-10C-MRATE Rate-Selectability Overview | 336](#)
[MPC10E-15C-MRATE Overview](#)
[MPC10E-15C-MRATE Rate-Selectability Overview | 332](#)

## Release History Table

| Release                | Description                                                                                                    |
|------------------------|----------------------------------------------------------------------------------------------------------------|
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE. |
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE. |

## RELATED DOCUMENTATION

[Introduction to Rate Selectability | 313](#)

[Configuring 100-Gigabit Ethernet MICs/PICs | 243](#)

## Configuring Rate Selectability

### IN THIS SECTION

- [Configuring Port Speed | 380](#)
- [Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers | 381](#)
- [Configuring Rate Selectability on MIC-MRATE to Enable Different Port Speeds | 385](#)
- [Configuring Rate Selectability on MPC7E \(Multi-Rate\) to Enable Different Port Speeds | 389](#)
- [Configuring Rate Selectability on MX10003 MPC to Enable Different Port Speeds | 394](#)
- [Configuring Rate Selectability on MX204 to Enable Different Port Speeds | 398](#)
- [Configuring Rate Selectability on JNP10K-2101 MPC to Enable Different Port Speeds | 402](#)
- [Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds | 406](#)
- [Configuring Rate Selectability on MPC10E-10C-MRATE to Enable Different Port Speeds | 410](#)
- [Configuring Rate Selectability on the MX2K-MPC11E to Enable Different Port Speeds | 414](#)
- [Configure Rate Selectability on ACX5448-D and ACX5448-M Routers | 417](#)
- [Channelize Interfaces on ACX5448-D and ACX5448-M Routers | 419](#)

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 ;
```

3. 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 Number | Speed  |
|-------------|--------|
| 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 68 on page 383](#) lists the guidelines to configure 4x100 Gbs on the PTX10003 routers in a tabular format.

**Table 68: 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 \(Ethernet\) | 1359](#)

[Unused | 1421](#)

## Configuring Rate Selectability on MIC-MRATE to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Rate Selectability on MIC-MRATE at MIC Level | 385](#)
- [Configuring Rate Selectability on MIC-MRATE at Port Level | 387](#)

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 Rate Selectability 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 Rate Selectability to Prevent Oversubscription”](#) on page 364.

```
[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 Rate Selectability 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 Rate Selectability” on page 318](#).

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.

SEE ALSO

|                                                        |
|--------------------------------------------------------|
| <i>number-of-ports</i>                                 |
| <i>pic-mode</i>                                        |
| <i>speed</i>                                           |
| <a href="#">Understanding Rate Selectability   314</a> |

### Configuring Rate Selectability on MPC7E (Multi-Rate) to Enable Different Port Speeds

IN THIS SECTION

- [Configuring Rate Selectability at PIC Level | 390](#)
- [Configuring Rate Selectability at Port Level | 391](#)

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 Rate Selectability 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 69 on page 391](#) lists the physical ports that are enabled when the **number-of-ports** statement is configured.

**Table 69: Active Physical Ports on MPC7E-MRATE MPC Based on the number-of-ports Configuration**

| Ports Configured<br>( <b>number-of-ports</b><br>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 Rate Selectability 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, you must reset the MPC for the configuration to take effect. Because resetting the MPC takes several minutes and as it affects all the Packet Forwarding Engines, you can choose to use the **request chassis pic pic-slot pic-slot-number fpc-slot fpc-slot-number (online | offline)** command to apply your configuration changes quickly. 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.

SEE ALSO

*MPC7E (Multi-Rate) on MX Series Routers Overview*

*pic-mode*

*speed*

*number-of-ports*

## Configuring Rate Selectability on MX10003 MPC to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Rate Selectability on MX10003 MPC at MIC/PIC Level | 394](#)
- [Configuring Rate Selectability on MX10003 MPC at Port Level | 396](#)

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 Rate Selectability 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 Rate Selectability to Prevent Oversubscription on MX10003 MPC” on page 367](#).

```
[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 Rate Selectability 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 Rate Selectability” on page 318](#).

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 Rate-Selectability Overview” on page 320](#) for more details.

## SEE ALSO

*number-of-ports*

*pic-mode*

*speed*

[Understanding Rate Selectability | 314](#)

## Configuring Rate Selectability on MX204 to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Rate Selectability on MX204 at PIC Level | 398](#)
- [Configuring Rate Selectability on MX204 at Port Level | 400](#)

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 Rate-Selectability Overview” on page 325](#).
- 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 Rate Selectability 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” on page 369](#).

```
[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 Rate Selectability 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 Rate Selectability” on page 318](#).

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 Rate-Selectability Overview” on page 320](#) for more details.

SEE ALSO

|                                                        |
|--------------------------------------------------------|
| <i>number-of-ports</i>                                 |
| <i>pic-mode</i>                                        |
| <i>speed</i>                                           |
| <a href="#">Understanding Rate Selectability   314</a> |
| <i>MX204 Router Overview</i>                           |
| <i>request chassis pic</i>                             |

Configuring Rate Selectability on JNP10K-2101 MPC to Enable Different Port Speeds

IN THIS SECTION

- [Configuring Rate Selectability on JNP10K-2101 MPC at PIC Level | 402](#)
- [Configuring Rate Selectability on JNP10K-LC2101 MPC at Port Level | 404](#)

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 Rate Selectability on JNP10K-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 JNP10K-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 Rate Selectability on JNP10K-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 JNP10K-LC2101, see [“Guidelines for Configuring Rate Selectability” on page 318](#).

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.

#### SEE ALSO

---

*number-of-ports*

---

*pic-mode*

---

*speed*

---

[Understanding Rate Selectability | 314](#)

## Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Rate Selectability on MPC10E-15C-MRATE at PIC Level | 406](#)
- [Configuring Rate Selectability on MPC10E-15C-MRATE at Port Level | 408](#)

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 in a PIC, configure rate selectability at the port level, in which case only the ports that are configured are created. To configure rate selectability at the port level, use the **speed** statement to specify the speed of individual ports.
- (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** is 2, then two IFDs are created, namely xe-x/y/z:0 and xe-x/y/z:1.

The **number-of-sub-ports** configuration statement can be used with rate selectability configuration at both PIC level and port level. The **number-of-sub-ports** configuration statement is effective only when the port speed is 10-Gbps. By default, four channels (or IFDs) is created for the 10-Gbps ports.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE.

This topic describes how to configure port speeds at the port level and at the PIC level.

### **Configuring Rate Selectability on MPC10E-15C-MRATE 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 MPC10E-15C-MRATE MPC supports port speeds of 10 Gbps, 40 Gbps, 100 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 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, 100G, and 400G.

```
[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 MPC10E-15C-MRATE MPC, see [“Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on MPC10E-15C-MRATE” on page 373](#).

```
[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 four ports on the MPC with port speed of 10 Gbps. The other ports are disabled.

### **Configuring Rate Selectability on MPC10E-15C-MRATE at Port Level**

To configure different port speeds for each port, you configure rate selectability at the port level. Only the configured ports 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. For guidelines on configuring rate selectability, see [“Guidelines for Configuring Rate Selectability” on page 318](#).

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, 100G, and 400G speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G | 40G | 100G | 400G)
```

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 two ports with port speed of 10 Gbps and one port with port speed of 40 Gbps.

SEE ALSO

---

*number-of-ports*

---

[number-of-sub-ports | 1291](#)

---

*pic-mode*

---

*speed*

---

[Understanding Rate Selectability | 314](#)

---

[MPC10E-15C-MRATE Rate-Selectability Overview | 332](#)

---

[Interface Naming Conventions for MPC10E-15C-MRATE MPC | 355](#)

## Configuring Rate Selectability on MPC10E-10C-MRATE to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Rate Selectability on MPC10E-10C-MRATE at PIC Level | 410](#)
- [Configuring Rate Selectability on MPC10E-10C-MRATE at Port Level | 412](#)

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 in a PIC, configure rate selectability at the port level, in which case only the ports that are configured are created. To configure rate selectability at the port level, use the **speed** statement to specify the speed of individual ports.
- (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** is 2, then two IFDs are created, namely xe-x/y/z:0 and xe-x/y/z:1.

The **number-of-sub-ports** configuration statement can be used with rate selectability configuration at both PIC level and port level. The **number-of-sub-ports** configuration statement is effective only when the port speed is 10-Gbps. By default, four channels (or IFDs) is created for the 10-Gbps ports.

Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE.

This topic describes how to configure port speeds at the port level and at the PIC level.

### **Configuring Rate Selectability on MPC10E-10C-MRATE 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 MPC10E-10C-MRATE MPC supports port speeds of 10 Gbps, 40 Gbps, 100 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 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, 100G, and 400G.

```
[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 MPC10E-10C-MRATE MPC, see [“Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on MPC10E-10C-MRATE” on page 376](#).

```
[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 four ports on the MPC with port speed of 10 Gbps. The other ports are disabled.

### **Configuring Rate Selectability on MPC10E-10C-MRATE at Port Level**

To configure different port speeds for each port, you configure rate selectability at the port level. Only the configured ports 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. For guidelines on configuring rate selectability, see [“Guidelines for Configuring Rate Selectability” on page 318](#).

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, 100G, and 400G speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10G | 40G | 100G | 400G)
```

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 two ports with port speed of 10 Gbps and one port with port speed of 40 Gbps.

SEE ALSO

*number-of-ports*

[number-of-sub-ports](#) | 1291

*pic-mode*

*speed*

[Understanding Rate Selectability](#) | 314

[MPC10E-10C-MRATE Rate-Selectability Overview](#) | 336

[Interface Naming Conventions for MPC10E-10C-MRATE MPC](#) | 356

*MPC10E-15C-MRATE Overview*

[MPC10E-15C-MRATE Rate-Selectability Overview](#) | 332

## Configuring Rate Selectability on the MX2K-MPC11E to Enable Different Port Speeds

### IN THIS SECTION

- [Configuring Rate Selectability on the MX2K-MPC11E at the PIC Level | 414](#)
- [Configuring Rate Selectability on the MX2K-MPC11E at the Port Level | 415](#)

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 in a PIC, configure rate selectability at the port level, in which case only the ports that are configured are created. 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 Rate Selectability on the MX2K-MPC11E at the 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.

**NOTE:** The MX2K-MPC11E supports PIC level configuration using **pic-mode** only for port speeds of 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 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. Supported value for the **pic-speed** option is 100G for MX2K-MPC11E.

**NOTE:** If you do not specify **pic-mode**, then by default the speed is 100-Gbps.

```
[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 100G
```

3. Verify the configuration.

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

4. Commit your configuration changes.

In this example, you have configured all five ports on the MPC with port speed of 100 Gbps.

#### **Configuring Rate Selectability on the MX2K-MPC11E at the Port Level**

Configure rate selectability at the port level if you want to operate individual ports of a PIC at different port speeds. Only the ports you configure are created. Other ports are automatically disabled.

**NOTE:** While configuring rate selectability, when you switch to PIC mode from port mode or vice-versa, the PIC is reset automatically. For guidelines on configuring rate selectability, see [“Guidelines for Configuring Rate Selectability” on page 318](#).

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. For MX2K-MPC11E, only Port 0 of every PIC, supports port speeds of 10 Gbps and 40 Gbps. If you configure the port speed for any other port, an alarm is generated indicating that the port speed configuration is invalid.

**NOTE:** By using the **speed** statement, you are specifying the speed of a single interface on that port.

```
[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 40G
user@host# set port 1 speed 100G
```

3. Verify the configuration.

```
[edit chassis fpc 0 pic 0]
user@host# show
port 0 {
    speed 40g;
}
port 1 {
    speed 100g;
}
```

4. Commit your configuration changes.

In this example, you have configured port 0 with port speed of 40 Gbps and port 1 with port speed of 100 Gbps.

**NOTE:** Only Port 0 of every PIC supports port speeds of 10 Gbps and 40 Gbps. The default port speed for all ports is 100 Gbps. You can also configure the port speed to be 100 Gbps.

#### SEE ALSO

*pic-mode*

*speed*

[Understanding Rate Selectability | 314](#)

### Configure Rate Selectability 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 Rate Selectability](#) | 313

## 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*

Release History Table

| Release                | Description                                                                                                    |
|------------------------|----------------------------------------------------------------------------------------------------------------|
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-15C-MRATE. |
| <a href="#">19.4R1</a> | Starting with Junos OS Release 19.4R1, you can now configure 400-Gbps speed on port 4 of the MPC10E-10C-MRATE. |
| <a href="#">15.1</a>   | Starting with Junos OS Release 15.1, some PICs support multiple port speeds.                                   |

RELATED DOCUMENTATION

|                                                                           |
|---------------------------------------------------------------------------|
| <a href="#">Introduction to Rate Selectability   313</a>                  |
| <a href="#">Interface Naming Conventions for Rate Selectability   345</a> |

# Configuring Optical Transport network

## IN THIS CHAPTER

- [Understanding Optical Transport Network \(OTN\) | 423](#)
- [Supported OTN and Optics Options | 452](#)
- [Forward Error Correction \(FEC\) and Bit Error Rate \(BER\) | 497](#)
- [Dense Wavelength Division Multiplexing \(DWDM\) Interface Wavelength | 505](#)
- [Configuring OTN | 508](#)
- [ODU Path Delay Measurement for Performance Monitoring | 537](#)

## Understanding Optical Transport Network (OTN)

### IN THIS SECTION

- [10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)
- [100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)
- [Understanding the P1-PTX-24-10G-W-SFPP PIC | 426](#)
- [Understanding the features of ACX6360 | 430](#)
- [Understanding the P2-100GE-OTN PIC | 433](#)
- [Understanding the MIC3-100G-DWDM MIC | 437](#)
- [Understanding the PTX-5-100G-WDM PIC | 441](#)
- [Understanding the PTX10K-LC1104 Line Card | 445](#)
- [Interface Mapping and Modulation format for ACX6360 | 447](#)
- [Interface Mapping and Modulation format for PTX10K-LC1104 Line Card | 450](#)

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. 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](#) | [1299](#)

## 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” on page 498](#) 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” on page 433](#) for more information.

See [“Supported OTN Options on PTX Series Routers” on page 453](#) 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-fpc/pic/port]** hierarchy level.

Use the **set optics-options** statement at the **[edit interfaces interfaceType-fpc/pic/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 | 512](#)

---

[show interfaces diagnostics optics \(Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port\) | 1580](#)

---

[optics-options | 1297](#)

---

[otn-options | 1299](#)

## Understanding the P1-PTX-24-10G-W-SFPP PIC

### IN THIS SECTION

- [Interface Features | 426](#)
- [Layer 2 and Layer 3 Features | 428](#)
- [OTN Alarms and Defects | 429](#)
- [TCA Alarms | 430](#)

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 gige-ether-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 | 508](#)

## **Understanding the features of ACX6360**

### **IN THIS SECTION**

- [Interface Features | 431](#)
- [Section | ?](#)
- [OTN Alarms and Defects | 431](#)
- [TCA Alarms | 432](#)

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 | 508](#)

## Understanding the P2-100GE-OTN PIC

### IN THIS SECTION

- [Interface Features | 434](#)
- [Layer 2 and Layer 3 Features | 435](#)
- [OTN Alarms and Defects | 436](#)
- [TCA Alarms | 437](#)

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” on page 86](#).

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 together-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 | 516](#)

## **Understanding the MIC3-100G-DWDM MIC**

### **IN THIS SECTION**

- [Interface Features | 438](#)
- [Layer 2 and Layer 3 Features | 439](#)
- [OTN Alarms and Defects | 439](#)

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)
  - HGFEC (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 | 521](#)

[Configuring Packet Optical Networks with PTX Series Devices](#)

## Understanding the PTX-5-100G-WDM PIC

### IN THIS SECTION

- [Interface Features | 441](#)
- [Layer 2 and Layer 3 Features | 443](#)
- [OTN Alarms and Defects | 443](#)

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 | 526](#)

*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 | 445](#)
- [OTN Alarms and Defects | 446](#)

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 | 532](#)

### 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<br>et-0/1/1     |
|                 | 16QAM-200G        | et-0/1/0<br>et-0/1/1     |
| 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-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
| ot-0/1/4        | QPSK-100G         | et-0/1/8                 |
|                 | 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   |
|                 | 16QAM-200G        | et-0/1/14<br>et-0/1/15   |

SEE ALSO

### 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     |
|                 | 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-" interface | Modulation Format | Mapped "et" interface(s) |
|-----------------|-------------------|--------------------------|
| 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     |
| 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 | 508](#)

Release History Table

| Release                | Description                                                                                                                                                                                                                   |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">15.1F5</a> | 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.                     |
| <a href="#">15.1</a>   | 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.                                 |
| <a href="#">15.1</a>   | 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

|                                                                               |
|-------------------------------------------------------------------------------|
| <a href="#">Configuring OTN   508</a>                                         |
| <a href="#">Forward Error Correction (FEC) and Bit Error Rate (BER)   497</a> |

Supported OTN and Optics Options

IN THIS SECTION

- Supported OTN Options on PTX Series Routers | [453](#)
- Supported OTN Options on MX Series Routers | [462](#)
- Supported OTN Options on ACX6360 Routers | [471](#)
- Supported OTN Options on ACX5448-D Routers | [477](#)
- Supported OTN Options on PTX10008 and PTX10016 Series Routers | [482](#)
- Supported Optics Options on ACX6360 and ACX5448-D Routers | [489](#)
- Supported Optics Options on PTX10008 and PTX10016 Series Routers | [493](#)

Use this topic for information about the supported optics options and otn options on specific devices.



## Supported OTN Options on PTX Series Routers

Table 70 on page 453 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 70: 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-2-100G-WDM<br>(PTX5000) |
|-----------------------------------------------------------------|-------------------------------------------------------|------------------------------------------|--------------------------------|--------------------------------|
| <code>bytes (otn-options)</code>                                | <code>transmit-payload-type value</code>              | 13.2/13.3                                | 14.1R2<br>14.2                 | 14.2                           |
| <code>fec</code>                                                | <code>(efec   gfec   gfec-sdfec   none   ufec)</code> | 13.2/13.3                                | 14.1R2<br>14.2<br>(gfec, none) | 14.2                           |
| <code>insert-odu-lck</code>                                     | -                                                     | 13.2/13.3                                | 14.1R2<br>14.2                 | 14.2                           |
| <code>insert-odu-oci</code>                                     | -                                                     | 13.2/13.3                                | 14.1R2<br>14.2                 | 14.2                           |
| <code>is-ma</code><br>  <code>no-is-ma</code>                   | -                                                     | 13.2/13.3                                | NA                             | 14.2                           |
| <code>laser-enable</code><br>  <code>no-laser-enable</code>     | -                                                     | 13.2/13.3                                | 14.1R2<br>14.2                 | 14.2                           |
| <code>line-loopback</code><br>  <code>no-line-loopback</code>   | -                                                     | 13.2/13.3                                | 14.1R2<br>14.2                 | 14.2                           |
| <code>local-loopback</code><br>  <code>no-local-loopback</code> | -                                                     | 13.2/13.3                                | 14.1R2<br>14.2                 | 14.2                           |

Table 70: 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-2-100G-WDM<br>(PTX5000) |
|----------------------------------------------------------|---------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
| odu-delay-management                                     | bypass  <br>no-bypass                       | 13.2/13.3                                | NA                        | NA                             |
|                                                          | monitor-end-point  <br>no-monitor-end-point | 13.2/13.3                                | NA                        | NA                             |
|                                                          | number-of-frames value                      | 13.2/13.3                                | NA                        | NA                             |
|                                                          | no-start-measurement  <br>start-measurement | 13.2/13.3                                | NA                        | NA                             |
| odu-signal-degrade                                       | ber-threshold-clear value                   | NA<br>NA                                 | 14.1R2<br>14.2            | NA                             |
|                                                          | ber-threshold-signal-degrade<br>value       | NA<br>NA                                 | 14.1R2<br>14.2            | NA                             |
|                                                          | interval value                              | NA                                       | 14.1R2<br>14.2            | NA                             |
|                                                          |                                             |                                          |                           |                                |
| odu-ttim-action-enable<br> <br>no-odu-ttim-action-enable | -                                           | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
| otu-ttim-action-enable<br> <br>no-otu-ttim-action-enable | -                                           | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
| prbs   no-prbs                                           | -                                           | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |

Table 70: 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-2-100G-WDM<br>(PTX5000) |
|-------------------------|--------------------------------------------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
| preemptive-fast-reroute | backward-frr-enable  <br>no-backward-frr-enable                                | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                         | signal-degrade-monitor-enable<br> <br>no-signal-degrade-monitor-enable         | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                         | odu-backward-frr-enable  <br>no-odu-backward-frr-enable                        | NA                                       | 14.1R2<br>14.2            | NA                             |
|                         | odu-signal-degrade-monitor-enable<br> <br>no-odu-signal-degrade-monitor-enable | NA                                       | 14.1R2<br>14.2            | NA                             |
| rate                    | fixed-stuff-bytes  <br>no-fixed-stuff-bytes                                    | 13.2/13.3                                | NA                        | 14.2                           |
|                         | oc192                                                                          | 13.2/13.3                                | NA                        | 14.2                           |
|                         | otu4                                                                           | 13.2/13.3                                | 14.1R2<br>14.2            | NA                             |
|                         | pass-through  <br>no-pass-through                                              | 13.2/13.3                                | NA                        | NA                             |
| signal-degrade          | ber-threshold-clear <b>value</b>                                               | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                         | ber-threshold-signal-degrade<br><b>value</b>                                   | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                         | interval <b>value</b>                                                          | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |

Table 70: 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-2-100G-WDM<br>(PTX5000) |
|-----------|---------------------------------------------------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
| tca       | odu-tca-bbe<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)    | 13.2/13.3                                | 14.2                      | 14.2                           |
|           | odu-tca-bbe-fe<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs) | 13.2/13.3                                | NA                        | 14.2                           |
|           | odu-tca-es<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)     | 13.2/13.3                                | 14.2                      | 14.2                           |
|           | odu-tca-es-fe<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)  | 13.2/13.3                                | NA                        | 14.2                           |
|           | odu-tca-ses<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)    | 13.2/13.3                                | 14.2                      | 14.2                           |
|           | odu-tca-ses-fe<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs) | 13.2/13.3                                | NA                        | 14.2                           |
|           | odu-tca-uas<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)    | 13.2/13.3                                | 14.2                      | 14.2                           |

Table 70: 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-2-100G-WDM<br>(PTX5000) |
|-----------|-----------------------------------------------------------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
|           | <b>otu-tca-bbe</b><br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)     | 13.2/13.3                                | 14.2                      | 14.2                           |
|           | <b>otu-tca-bbe-fe</b><br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)  | 13.2/13.3                                | NA                        | 14.2                           |
|           | <b>otu-tca-es</b><br>(enable-tca  <br>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  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)   | 13.2/13.3                                | NA                        | 14.2                           |
|           | <b>otu-tca-fec-ber</b><br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs) | 13.2/13.3                                | 14.2                      | 14.2                           |
|           | <b>otu-tca-ses</b><br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)     | 13.2/13.3                                | 14.2                      | 14.2                           |
|           | <b>otu-tca-ses-fe</b><br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)  | 13.2/13.3                                | NA                        | 14.2                           |

Table 70: 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-2-100G-WDM<br>(PTX5000) |
|--------------------------------------|---------------------------------------------------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
|                                      | otu-tca-uas<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs)    | 13.2/13.3                                | 14.2                      | 14.2                           |
|                                      | otu-tca-uas-fe<br>(enable-tca  <br>no-enable-tca  <br>threshold  <br>threshold-24hrs) | 13.2/13.3                                | NA                        | 14.2                           |
| <a href="#">transport-monitoring</a> | -                                                                                     | NA                                       | NA                        | 14.2                           |

Table 70: 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-2-100G-WDM<br>(PTX5000) |
|------------------------------------------------------|-----------------------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
| <a href="#">trigger</a><br><b>trigger-identifier</b> | oc-lof<br>(hold-time (down   up)  <br>ignore)             | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | oc-lom<br>(hold-time (down   up)  <br>ignore)             | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | oc-los<br>(hold-time (down   up)  <br>ignore)             | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | oc-tsf<br>(hold-time (down   up)  <br>ignore)             | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | oc-wavelength-lock<br>(hold-time (down   up)  <br>ignore) | 13.2/13.3                                | NA                        | 14.2                           |
|                                                      | odu-ais<br>(hold-time (down   up)  <br>ignore)            | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | odu-bdi<br>(hold-time (down   up)  <br>ignore)            | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | odu-bei<br>(hold-time (down   up)  <br>ignore)            | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | odu-iae<br>(hold-time (down   up)  <br>ignore)            | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      | odu-lck<br>(hold-time (down   up)  <br>ignore)            | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                                                      |                                                           | 13.2/13.3                                |                           | 14.2                           |

Table 70: 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-2-100G-WDM<br>(PTX5000) |
|-----------|--------------------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
|           | <b>odu-oci</b><br>(hold-time (down   up)   ignore)     |                                          | 14.1R2<br>14.2            |                                |
|           | <b>odu-sd</b> (hold-time (down   up)   ignore)         | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>odu-ttim</b>                                        | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>opu-ptim</b><br>(hold-time (down   up)   ignore)    | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>otu-ais</b><br>(hold-time (down   up)   ignore)     | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>otu-bdi</b><br>(hold-time (down   up)   ignore)     | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>otu-fec-deg</b><br>(hold-time (down   up)   ignore) | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>otu-fec-exe</b><br>(hold-time (down   up)   ignore) | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>otu-iae</b><br>(hold-time (down   up)   ignore)     | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           | <b>otu-sd</b><br>(hold-time (down   up)   ignore)      | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|           |                                                        | 13.2/13.3                                |                           | 14.2                           |



Table 70: 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-2-100G-WDM<br>(PTX5000) |
|----------------|----------------------------------------------|------------------------------------------|---------------------------|--------------------------------|
|                | otu-ttim (hold-time<br>(down   up)   ignore) |                                          | 14.1R2<br>14.2            |                                |
| tti-identifier | odu-dapi<br>identifier                       | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                | odu-expected-<br>receive-dapi<br>identifier  | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                | odu-expected-<br>receive-sapi<br>identifier  | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                | odu-sapi<br>identifier                       | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                | otu-dapi<br>identifier                       | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                | otu-expected-<br>receive-dapi<br>identifier  | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                | otu-expected-<br>receive-sapi<br>identifier  | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |
|                | otu-sapi<br>identifier                       | 13.2/13.3                                | 14.1R2<br>14.2            | 14.2                           |

SEE ALSO

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## Supported OTN Options on MX Series Routers

Table 71 on page 462 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 71: Statements Supported on 100-Gigabit Ethernet MICs on MX Series Routers**

| Statement                                                       | Options                                                                   | MIC6-100G-CFP2<br>(MX2010 / MX2020) | MIC3-100G-DWDM<br>(MX240, MX480, MX960,<br>MX2010, and MX2020) |
|-----------------------------------------------------------------|---------------------------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
| <code>bytes</code>                                              | <code>transmit-payload-type value</code>                                  | NA                                  | 15.1F5                                                         |
| <code>fec</code>                                                | <code>(efec   gfec   gfec-sdfec   hgfec  <br/>none   dsfec   ufec)</code> | 13.3R3<br><br>(ufec)                | 15.1F5<br><br>(gfec,hgfec,sdfec)                               |
| <code>insert-<br/>odu-lck</code>                                | -                                                                         | 13.3R3                              | 15.1F5                                                         |
| <code>insert-<br/>odu-oci</code>                                | -                                                                         | 13.3R3                              | 15.1F5                                                         |
| <code>is-ma</code><br>  <code>no-is-ma</code>                   | -                                                                         | 13.3R3                              | 15.1F5                                                         |
| <code>laser-enable</code><br>  <code>no-laser-enable</code>     | -                                                                         | 13.3R3                              | 15.1F5                                                         |
| <code>line-loopback</code><br>  <code>no-line-loopback</code>   | -                                                                         | 13.3R3                              | 15.1F5                                                         |
| <code>local-loopback</code><br>  <code>no-local-loopback</code> | -                                                                         | 13.3R3                              | 15.1F5                                                         |

Table 71: 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) |
|----------------------------------------------------------|---------------------------------------------|-------------------------------------|----------------------------------------------------------------|
| odu-delay-management                                     | bypass  <br>no-bypass                       | NA                                  | 15.1F5                                                         |
|                                                          | monitor-end-point  <br>no-monitor-end-point | NA                                  | 15.1F5                                                         |
|                                                          | number-of-frames <i>value</i>               | NA                                  | 15.1F5                                                         |
|                                                          | no-start-measurement  <br>start-measurement | NA                                  | 15.1F5                                                         |
| signal-degrade                                           | ber-threshold-clear <i>value</i>            | 13.3R3                              | 15.1F5                                                         |
|                                                          | ber-threshold-signal-degrade <i>value</i>   | 13.3R3                              | 15.1F5                                                         |
|                                                          | interval <i>value</i>                       | 13.3R3                              | 15.1F5                                                         |
| odu-ttim-action-enable<br> <br>no-odu-ttim-action-enable | -                                           | 13.3R3                              | 15.1F5                                                         |
| otu-ttim-action-enable<br> <br>no-otu-ttim-action-enable | -                                           | 13.3R3                              | 15.1F5                                                         |
| prbs   no-prbs                                           | -                                           | 13.3R3                              | 15.1F5                                                         |

Table 71: 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) |
|-------------------------|--------------------------------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
| preemptive-fast-reroute | backward-frr-enable  <br>no-backward-frr-enable                                | 13.3R3                              | 15.1F5                                                         |
|                         | signal-degrade-monitor-enable  <br>no-signal-degrade-monitor-enable            | 13.3R3                              | 15.1F5                                                         |
|                         | odu-backward-frr-enable  <br>no-odu-backward-frr-enable                        | NA                                  | 15.1F5                                                         |
|                         | odu-signal-degrade-monitor-enable<br> <br>no-odu-signal-degrade-monitor-enable | NA                                  | 15.1F5                                                         |
| rate                    | fixed-stuff-bytes  <br>no-fixed-stuff-bytes                                    | 13.3R3                              | 15.1F5                                                         |
|                         | oc192                                                                          | 13.3R3                              | 15.1F5                                                         |
|                         | otu4                                                                           | 13.3R3                              | 15.1F5                                                         |
|                         | pass-through                                                                   | 13.3R3<br>(pass-through)            | 15.1F5<br>(pass-through)                                       |
| signal-degrade          | ber-threshold-clear <b>value</b>                                               | 13.3R3                              | 15.1F5                                                         |
|                         | ber-threshold-signal-degrade <b>value</b>                                      | 13.3R3                              | 15.1F5                                                         |
|                         | interval <b>value</b>                                                          | 13.3R3                              | 15.1F5                                                         |
|                         | q-threshold-signal-degrade                                                     | NA                                  | 15.1F5                                                         |
|                         | q-threshold-signal-degrade-clear                                               | NA                                  | 15.1F5                                                         |

Table 71: 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) |
|-----------|---------------------------------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
| tca       | odu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 13.3R3                              | 15.1F5                                                         |
|           | otu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 13.3R3                              | 15.1F5                                                         |
|           | otu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 13.3R3                              | 15.1F5                                                         |
|           | otu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 13.3R3                              | 15.1F5                                                         |

Table 71: 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-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)   | 13.3R3                              | 15.1F5                                                         |
|                      | otu-tca-fec-ber<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 13.3R3                              | 15.1F5                                                         |
|                      | otu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 13.3R3                              | 15.1F5                                                         |
|                      | otu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 13.3R3                              | 15.1F5                                                         |
|                      | otu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 13.3R3                              | 15.1F5                                                         |
|                      | otu-tca-uas-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 13.3R3                              | 15.1F5                                                         |
| transport-monitoring | -                                                                                | NA                                  | 15.1F5                                                         |

Table 71: 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) |
|-------------------------------|-----------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
| trigger<br>trigger-identifier | oc-lof<br>(hold-time (down   up)  <br>ignore)             | 13.3R3                              | 15.1F5                                                         |
|                               | oc-lom<br>(hold-time (down   up)  <br>ignore)             | 13.3R3                              | 15.1F5                                                         |
|                               | oc-los<br>(hold-time (down   up)  <br>ignore)             | 13.3R3                              | 15.1F5                                                         |
|                               | oc-tsfc<br>(hold-time (down   up)  <br>ignore)            | NA                                  | 15.1F5                                                         |
|                               | oc-wavelength-lock<br>(hold-time (down   up)  <br>ignore) | NA                                  | 15.1F5                                                         |
|                               | odu-ais<br>(hold-time (down   up)  <br>ignore)            | 13.3R3                              | 15.1F5                                                         |
|                               | odu-bdi<br>(hold-time (down   up)  <br>ignore)            | 13.3R3                              | 15.1F5                                                         |
|                               | odu-bei<br>(hold-time (down   up)  <br>ignore)            | 13.3R3                              | 15.1F5                                                         |
|                               | odu-iae<br>(hold-time (down   up)  <br>ignore)            | 13.3R3                              | 15.1F5                                                         |
|                               | odu-lck<br>(hold-time (down   up)  <br>ignore)            | 13.3R3                              | 15.1F5                                                         |

Table 71: 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) |
|-----------|----------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
|           | odu-oci<br>(hold-time (down   up)  <br>ignore)     | 13.3R3                              | 15.1F5                                                         |
|           | odu-sd (hold-time<br>(down   up)   ignore)         | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-es                                         | 13.3R3                              | NA                                                             |
|           | odu-tca-ses                                        | 13.3R3                              | NA                                                             |
|           | odu-tca-uas                                        | 13.3R3                              | NA                                                             |
|           | odu-ttim                                           | 13.3R3                              | 15.1F5                                                         |
|           | opu-ptim<br>(hold-time (down   up)  <br>ignore)    | 13.3R3                              | 15.1F5                                                         |
|           | otu-ais<br>(hold-time (down   up)  <br>ignore)     | 13.3R3                              | 15.1F5                                                         |
|           | otu-bdi<br>(hold-time (down   up)  <br>ignore)     | 13.3R3                              | 15.1F5                                                         |
|           | otu-fec-deg<br>(hold-time (down   up)  <br>ignore) | NA                                  | 15.1F5                                                         |
|           | otu-fec-exe<br>(hold-time (down   up)  <br>ignore) | NA                                  | 15.1F5                                                         |
|           | otu-iae<br>(hold-time (down   up)  <br>ignore)     | 13.3R3                              | 15.1F5                                                         |



Table 71: 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-sd<br>(hold-time (down   up)  <br>ignore) | 13.3R3                              | 15.1F5                                                         |
|           | odu-tca-es                                    | 13.3R3                              | NA                                                             |
|           | odu-tca-ses                                   | 13.3R3                              | NA                                                             |
|           | odu-tca-uas                                   | 13.3R3                              | NA                                                             |
|           | otu-ttim (hold-time (down  <br>up)   ignore)  | 13.3R3                              | 15.1F5                                                         |

Table 71: 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) |
|--------------------------------|-------------------------------------------------------------------------------------------|-------------------------------------|----------------------------------------------------------------|
| <a href="#">tti-identifier</a> | odu-dapi<br><i>identifiery</i>                                                            | 13.3R3                              | 15.1F5                                                         |
|                                | odu-dapi-first-byte-nul  <br>no-odu-dapi-first-byte-nul                                   | NA                                  | 15.1F5                                                         |
|                                | odu-expected-receive-dapi<br><i>identifier</i>                                            | 13.3R3                              | 15.1F5                                                         |
|                                | odu-expected-receive-dapi-first-byte-nul  <br>no-odu-expected-receive-dapi-first-byte-nul | NA                                  | 15.1F5                                                         |
|                                | odu-expected-receive-sapi<br><i>identifier</i>                                            | 13.3R3                              | 15.1F5                                                         |
|                                | odu-sapi<br><i>identifiery</i>                                                            | 13.3R3                              | 15.1F5                                                         |
|                                | odu-sapi-first-byte-nul  <br>no-odu-sapi-first-byte-nul                                   | NA                                  | 15.1F5                                                         |
|                                | otu-dapi<br><i>identifier</i>                                                             | 13.3R3                              | 15.1F5                                                         |
|                                | otu-dapi-first-byte-nul  <br>no-otu-dapi-first-byte-nul                                   | NA                                  | 15.1F5                                                         |
|                                | otu-expected-<br>receive-dapi<br><i>identifier</i>                                        | 13.3R3                              | 15.1F5                                                         |
|                                | otu-expected-receive-dapi-first-byte-nul  <br>no-otu-expected-receive-dapi-first-byte-nul | NA                                  | 15.1F5                                                         |
|                                | otu-expected-<br>receive-sapi<br><i>identifier</i>                                        | 13.3R3                              | 15.1F5                                                         |
|                                | otu-expected-receive-sapi-first-byte-nul  <br>no-otu-expected-receive-sapi-first-byte-nul | NA                                  | 15.1F5                                                         |

Table 71: 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>otu-sapi<br/>identifier</b> | 13.3R3                              | 15.1F5                                                         |
|           | <b>otu-sapi-first-byte-nul</b> | NA                                  | 15.1F5                                                         |

SEE ALSO

[Configuring OTN | 508](#)

## Supported OTN Options on ACX6360 Routers

Table 72 on page 471 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 72: Statements Supported on ACX6360 Routers

| Statement                                                    | Options                                 | Junos Version |
|--------------------------------------------------------------|-----------------------------------------|---------------|
| <a href="#">bytes (otn-options)</a>                          | <b>transmit-payload-<br/>type value</b> | 18.3R1        |
| <b>insert-<br/>odu-lck</b>                                   | -                                       | 18.3R1        |
| <b>insert-<br/>odu-oci</b>                                   | -                                       | 18.3R1        |
| <a href="#">is-ma</a><br>  <b>no-is-ma</b>                   | -                                       | 18.3R1        |
| <a href="#">line-loopback</a><br>  <b>no-line-loopback</b>   | -                                       | 18.3R1        |
| <a href="#">local-loopback</a><br>  <b>no-local-loopback</b> | -                                       | 18.3R1        |

Table 72: Statements Supported on ACX6360 Routers (continued)

| Statement                                                                       | Options                                                                                               | Junos Version |
|---------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|---------------|
| <code>odu-ttim-action-enable</code>  <br><code>no-odu-ttim-action-enable</code> | -                                                                                                     | 18.3R1        |
| <code>otu-ttim-action-enable</code>  <br><code>no-otu-ttim-action-enable</code> | -                                                                                                     | 18.3R1        |
| <code>prbs</code>   <code>no-prbs</code>                                        | -                                                                                                     | 18.3R1        |
| <code>preemptive-fast-reroute</code>                                            | <code>backward-frr-enable</code>   <code>no-backward-frr-enable</code>                                | 18.3R1        |
|                                                                                 | <code>signal-degrade-monitor-enable</code>  <br><code>no-signal-degrade-monitor-enable</code>         | 18.3R1        |
|                                                                                 | <code>odu-backward-frr-enable</code>   <code>no-odu-backward-frr-enable</code>                        | NA            |
|                                                                                 | <code>odu-signal-degrade-monitor-enable</code>  <br><code>no-odu-signal-degrade-monitor-enable</code> | NA            |

Table 72: Statements Supported on ACX6360 Routers (*continued*)

| Statement | Options                                                                         | Junos Version |
|-----------|---------------------------------------------------------------------------------|---------------|
| tca       | odu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1        |
|           | odu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1        |
|           | odu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.3R1        |
|           | odu-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 18.3R1        |
|           | odu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1        |
|           | odu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1        |
|           | odu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1        |
|           | otu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1        |
|           | otu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1        |
|           | otu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.3R1        |
|           |                                                                                 | 18.3R1        |

Table 72: Statements Supported on ACX6360 Routers (*continued*)

| Statement | Options                                                                         | Junos Version |
|-----------|---------------------------------------------------------------------------------|---------------|
|           | otu-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  |               |
|           | otu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1        |
|           | otu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1        |
|           | otu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1        |
|           | otu-tca-uas-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1        |

Table 72: Statements Supported on ACX6360 Routers (*continued*)

| Statement                               | Options                                                   | Junos Version |
|-----------------------------------------|-----------------------------------------------------------|---------------|
| <code>trigger trigger-identifier</code> | <code>oc-lof</code><br>(hold-time (down   up)   ignore)   | 18.3R1        |
|                                         | <code>oc-lom</code><br>(hold-time (down   up)   ignore)   | 18.3R1        |
|                                         | <code>oc-tsf</code><br>(hold-time (down   up)   ignore)   | 18.3R1        |
|                                         | <code>odu-ais</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |
|                                         | <code>odu-bdi</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |
|                                         | <code>odu-bei</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |
|                                         | <code>odu-iae</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |
|                                         | <code>odu-lck</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |
|                                         | <code>odu-oci</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |
|                                         | <code>odu-sd</code> (hold-time<br>(down   up)   ignore)   | 18.3R1        |
|                                         | <code>odu-ttim</code>                                     | 18.3R1        |
|                                         | <code>opu-ptim</code><br>(hold-time (down   up)   ignore) | 18.3R1        |
|                                         | <code>otu-ais</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |
|                                         | <code>otu-bdi</code><br>(hold-time (down   up)   ignore)  | 18.3R1        |

Table 72: Statements Supported on ACX6360 Routers (*continued*)

| Statement                 | Options                                            | Junos Version |
|---------------------------|----------------------------------------------------|---------------|
|                           | otu-iae<br>(hold-time (down   up)   ignore)        | 18.3R1        |
|                           | otu-sd<br>(hold-time (down   up)   ignore)         | 18.3R1        |
|                           | otu-ttim (hold-time (down   up)   ignore)          | 18.3R1        |
| tti <i>tti-identifier</i> | odu-dapi<br><i>identifier</i>                      | 18.3R1        |
|                           | odu-expected-<br>receive-dapi<br><i>identifier</i> | 18.3R1        |
|                           | odu-expected-<br>receive-sapi<br><i>identifier</i> | 18.3R1        |
|                           | odu-sapi<br><i>identifier</i>                      | 18.3R1        |
|                           | otu-dapi<br><i>identifier</i>                      | 18.3R1        |
|                           | otu-expected-<br>receive-dapi<br><i>identifier</i> | 18.3R1        |
|                           | otu-expected-<br>receive-sapi<br><i>identifier</i> | 18.3R1        |
|                           | otu-sapi<br><i>identifier</i>                      | 18.3R1        |

SEE ALSO

| [Configuring OTN](#) | 508



## Supported OTN Options on ACX5448-D Routers

Table 73 on page 477 lists the statements that are supported on ACX5448-D routers at the **[edit interfaces interface-name otn-options]** hierarchy level.

Table 73: Statements Supported on ACX5448-D Routers

| Statement                                                                       | Options                                                                                               | Junos OS |
|---------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|----------|
| <code>bytes (otn-options)</code>                                                | <code>transmit-payload-type value</code>                                                              | 19.2R1   |
| <code>insert-odu-lck</code>                                                     | -                                                                                                     | 19.2R1   |
| <code>insert-odu-oci</code>                                                     | -                                                                                                     | 19.2R1   |
| <code>is-ma</code><br>  <code>no-is-ma</code>                                   | -                                                                                                     | 19.2R1   |
| <code>line-loopback</code><br>  <code>no-line-loopback</code>                   | -                                                                                                     | 19.2R1   |
| <code>local-loopback</code><br>  <code>no-local-loopback</code>                 | -                                                                                                     | 19.2R1   |
| <code>odu-ttim-action-enable</code>  <br><code>no-odu-ttim-action-enable</code> | -                                                                                                     | 19.2R1   |
| <code>otu-ttim-action-enable</code>  <br><code>no-otu-ttim-action-enable</code> | -                                                                                                     | 19.2R1   |
| <code>prbs</code>   <code>no-prbs</code>                                        | -                                                                                                     | 19.2R1   |
| <code>preemptive-fast-reroute</code>                                            | <code>backward-frr-enable</code>   <code>no-backward-frr-enable</code>                                | 19.2R1   |
|                                                                                 | <code>signal-degrade-monitor-enable</code>  <br><code>no-signal-degrade-monitor-enable</code>         | 19.2R1   |
|                                                                                 | <code>odu-backward-frr-enable</code>   <code>no-odu-backward-frr-enable</code>                        | NA       |
|                                                                                 | <code>odu-signal-degrade-monitor-enable</code>  <br><code>no-odu-signal-degrade-monitor-enable</code> | NA       |

Table 73: Statements Supported on ACX5448-D Routers (*continued*)

| Statement | Options                                                                         | Junos OS |
|-----------|---------------------------------------------------------------------------------|----------|
| tca       | odu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 19.2R1   |
|           | odu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 19.2R1   |
|           | odu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 19.2R1   |
|           | odu-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 19.2R1   |
|           | odu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 19.2R1   |
|           | odu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 19.2R1   |
|           | odu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 19.2R1   |
|           | otu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 19.2R1   |
|           | otu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 19.2R1   |
|           | otu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 19.2R1   |
|           |                                                                                 | 19.2R1   |

Table 73: Statements Supported on ACX5448-D Routers (*continued*)

| Statement | Options                                                                         | Junos OS |
|-----------|---------------------------------------------------------------------------------|----------|
|           | otu-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  |          |
|           | otu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 19.2R1   |
|           | otu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 19.2R1   |
|           | otu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 19.2R1   |
|           | otu-tca-uas-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 19.2R1   |

Table 73: Statements Supported on ACX5448-D Routers (*continued*)

| Statement                               | Options                                                   | Junos OS |
|-----------------------------------------|-----------------------------------------------------------|----------|
| <code>trigger trigger-identifier</code> | <code>oc-lof</code><br>(hold-time (down   up)   ignore)   | 19.2R1   |
|                                         | <code>oc-lom</code><br>(hold-time (down   up)   ignore)   | 19.2R1   |
|                                         | <code>oc-tsf</code><br>(hold-time (down   up)   ignore)   | 19.2R1   |
|                                         | <code>odu-ais</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |
|                                         | <code>odu-bdi</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |
|                                         | <code>odu-bei</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |
|                                         | <code>odu-iae</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |
|                                         | <code>odu-lck</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |
|                                         | <code>odu-oci</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |
|                                         | <code>odu-sd</code> (hold-time<br>(down   up)   ignore)   | 19.2R1   |
|                                         | <code>odu-ttim</code>                                     | 19.2R1   |
|                                         | <code>opu-ptim</code><br>(hold-time (down   up)   ignore) | 19.2R1   |
|                                         | <code>otu-ais</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |
|                                         | <code>otu-bdi</code><br>(hold-time (down   up)   ignore)  | 19.2R1   |

Table 73: Statements Supported on ACX5448-D Routers (*continued*)

| Statement                 | Options                                            | Junos OS |
|---------------------------|----------------------------------------------------|----------|
|                           | otu-iae<br>(hold-time (down   up)   ignore)        | 19.2R1   |
|                           | otu-sd<br>(hold-time (down   up)   ignore)         | 19.2R1   |
|                           | otu-ttim (hold-time (down   up)   ignore)          | 19.2R1   |
| tti <i>tti-identifier</i> | odu-dapi<br><i>identifier</i>                      | 19.2R1   |
|                           | odu-expected-<br>receive-dapi<br><i>identifier</i> | 19.2R1   |
|                           | odu-expected-<br>receive-sapi<br><i>identifier</i> | 19.2R1   |
|                           | odu-sapi<br><i>identifier</i>                      | 19.2R11  |
|                           | otu-dapi<br><i>identifier</i>                      | 19.2R1   |
|                           | otu-expected-<br>receive-dapi<br><i>identifier</i> | 19.2R1   |
|                           | otu-expected-<br>receive-sapi<br><i>identifier</i> | 19.2R1   |
|                           | otu-sapi<br><i>identifier</i>                      | 19.2R1   |

SEE ALSO

| [Configuring OTN](#) | 508

## Supported OTN Options on PTX10008 and PTX10016 Series Routers

Table 74 on page 482 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 74: Statements Supported on PTX10K-LC1104 line cards

| Statement                                                       | Options                                               | Release | Interfaces Supported (ot/et) |
|-----------------------------------------------------------------|-------------------------------------------------------|---------|------------------------------|
| <code>bytes (otn-options)</code>                                | <code>transmit-payload-type value</code>              | 18.3R1  | et                           |
| <code>fec</code>                                                | <code>(efec   gfec   gfec-sdfec   none   ufec)</code> | 18.3R1  | ot                           |
| <code>insert-odu-lck</code>                                     | -                                                     | 18.3R1  | et                           |
| <code>insert-odu-oci</code>                                     | -                                                     | 18.3R1  | et                           |
| <code>is-ma</code><br>  <code>no-is-ma</code>                   | -                                                     | 18.3R1  | et                           |
| <code>laser-enable</code><br>  <code>no-laser-enable</code>     | -                                                     | 18.3R1  | ot                           |
| <code>line-loopback</code><br>  <code>no-line-loopback</code>   | -                                                     | 18.3R1  | ot                           |
| <code>local-loopback</code><br>  <code>no-local-loopback</code> | -                                                     | 18.3R1  | et                           |
| <code>modulation-format</code>                                  | <code>qpsk   8qam   16qam</code>                      | 18.3R1  | ot                           |

Table 74: Statements Supported on PTX10K-LC1104 line cards (continued)

| Statement                                             | Options                                                                     | Release | Interfaces Supported (ot/et) |
|-------------------------------------------------------|-----------------------------------------------------------------------------|---------|------------------------------|
| odu-delay-management                                  | bypass  <br>no-bypass                                                       | 18.3R1  | et                           |
|                                                       | monitor-end-point  <br>no-monitor-end-point                                 | 18.3R1  |                              |
|                                                       | number-of-frames value                                                      | 18.3R1  |                              |
|                                                       | no-start-measurement  <br>start-measurement                                 | 18.3R1  |                              |
| odu-signal-degrade                                    | ber-threshold-clear value                                                   | 18.3R1  | et                           |
|                                                       | ber-threshold-signal-degrade value                                          | 18.3R1  |                              |
|                                                       | interval value                                                              | 18.3R1  |                              |
| odu-ttim-action-enable  <br>no-odu-ttim-action-enable | -                                                                           | 18.3R1  | et                           |
| otu-ttim-action-enable  <br>no-otu-ttim-action-enable | -                                                                           | 18.3R1  | et                           |
| prbs   no-prbs                                        | -                                                                           | 18.3R1  | ot                           |
| preemptive-fast-reroute                               | backward-frr-enable  <br>no-backward-frr-enable                             | 18.3R1  | ot and et                    |
|                                                       | signal-degrade-monitor-enable  <br>no-signal-degrade-monitor-enable         | 18.3R1  |                              |
|                                                       | odu-backward-frr-enable  <br>no-odu-backward-frr-enable                     | 18.3R1  |                              |
|                                                       | odu-signal-degrade-monitor-enable  <br>no-odu-signal-degrade-monitor-enable | 18.3R1  |                              |

Table 74: Statements Supported on PTX10K-LC1104 line cards (*continued*)

| Statement      | Options                                   | Release | Interfaces Supported (ot/et) |
|----------------|-------------------------------------------|---------|------------------------------|
| rate           | fixed-stuff-bytes   no-fixed-stuff-bytes  | 18.3R1  | et                           |
|                | oc192                                     | 18.3R1  |                              |
|                | otu4                                      | 18.3R1  |                              |
|                | pass-through   no-pass-through            | 18.3R1  |                              |
| signal-degrade | ber-threshold-clear <b>value</b>          | 18.3R1  | ot and et                    |
|                | ber-threshold-signal-degrade <b>value</b> | 18.3R1  |                              |
|                | interval <b>value</b>                     | 18.3R1  |                              |



Table 74: Statements Supported on PTX10K-LC1104 line cards (*continued*)

| Statement | Options                                                                      | Release | Interfaces Supported (ot/et) |
|-----------|------------------------------------------------------------------------------|---------|------------------------------|
| tca       | odu-tca-bbe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  | ot and et                    |
|           | odu-tca-bbe-fe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                              |
|           | odu-tca-es<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                              |
|           | odu-tca-es-fe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1  |                              |
|           | odu-tca-ses<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                              |
|           | odu-tca-ses-fe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                              |
|           | odu-tca-uas<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                              |
|           | otu-tca-bbe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)    | 18.3R1  |                              |
|           | otu-tca-bbe-fe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                              |
|           | otu-tca-es<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                              |
|           |                                                                              | 18.3R1  |                              |

Table 74: Statements Supported on PTX10K-LC1104 line cards (*continued*)

| Statement                            | Options                                                                       | Release | Interfaces Supported (ot/et) |
|--------------------------------------|-------------------------------------------------------------------------------|---------|------------------------------|
|                                      | otu-tca-es-fe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)   |         |                              |
|                                      | otu-tca-fec-ber<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  |                              |
|                                      | otu-tca-ses<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                              |
|                                      | otu-tca-ses-fe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1  |                              |
|                                      | otu-tca-uas<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)     | 18.3R1  |                              |
|                                      | otu-tca-uas-fe<br>(enable-tca   no-enable-tca   threshold   threshold-24hrs)  | 18.3R1  |                              |
| <a href="#">transport-monitoring</a> | -                                                                             | 18.3R1  | et                           |

Table 74: Statements Supported on PTX10K-LC1104 line cards (*continued*)

| Statement                               | Options                                                             | Release | Interfaces Supported (ot/et) |
|-----------------------------------------|---------------------------------------------------------------------|---------|------------------------------|
| <code>trigger trigger-identifier</code> | <code>oc-lof</code><br>(hold-time (down   up)   ignore)             | 18.3R1  | ot and et                    |
|                                         | <code>oc-lom</code><br>(hold-time (down   up)   ignore)             | 18.3R1  |                              |
|                                         | <code>oc-los</code><br>(hold-time (down   up)   ignore)             | 18.3R1  |                              |
|                                         | <code>oc-tsf</code><br>(hold-time (down   up)   ignore)             | 18.3R1  |                              |
|                                         | <code>oc-wavelength-lock</code><br>(hold-time (down   up)   ignore) | 18.3R1  |                              |
|                                         | <code>odu-ais</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                              |
|                                         | <code>odu-bdi</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                              |
|                                         | <code>odu-bei</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                              |
|                                         | <code>odu-iae</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                              |
|                                         | <code>odu-lck</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                              |
|                                         | <code>odu-oci</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                              |
|                                         | <code>odu-sd</code> (hold-time (down   up)   ignore)                | 18.3R1  |                              |
|                                         | <code>odu-ttim</code>                                               | 18.3R1  |                              |
|                                         | <code>opu-ptim</code><br>(hold-time (down   up)   ignore)           | 18.3R1  |                              |

Table 74: Statements Supported on PTX10K-LC1104 line cards (*continued*)

| Statement | Options                                                | Release | Interfaces Supported<br>(ot/et) |
|-----------|--------------------------------------------------------|---------|---------------------------------|
|           | <b>otu-ais</b><br>(hold-time (down   up)   ignore)     | 18.3R1  |                                 |
|           | <b>otu-bdi</b><br>(hold-time (down   up)   ignore)     | 18.3R1  |                                 |
|           | <b>otu-fec-deg</b><br>(hold-time (down   up)   ignore) | 18.3R1  |                                 |
|           | <b>otu-fec-exe</b><br>(hold-time (down   up)   ignore) | 18.3R1  |                                 |
|           | <b>otu-iae</b><br>(hold-time (down   up)   ignore)     | 18.3R1  |                                 |
|           | <b>otu-sd</b><br>(hold-time (down   up)   ignore)      | 18.3R1  |                                 |
|           | <b>otu-ttim</b> (hold-time (down   up)<br>  ignore)    | 18.3R1  |                                 |

Table 74: Statements Supported on PTX10K-LC1104 line cards (*continued*)

| Statement      | Options                              | Release | Interfaces Supported (ot/et) |
|----------------|--------------------------------------|---------|------------------------------|
| tti-identifier | odu-dapi identifier                  | 18.3R1  | et                           |
|                | odu-expected-receive-dapi identifier | 18.3R1  |                              |
|                | odu-expected-receive-sapi identifier | 18.3R1  |                              |
|                | odu-sapi identifier                  | 18.3R1  |                              |
|                | otu-dapi identifier                  | 18.3R1  |                              |
|                | otu-expected-receive-dapi identifier | 18.3R1  |                              |
|                | otu-expected-receive-sapi identifier | 18.3R1  |                              |
|                | otu-sapi identifier                  | 18.3R1  |                              |

### Supported Optics Options on ACX6360 and ACX5448-D Routers

Table 75 on page 489 lists the statements that are supported on ACX6360 and ACX5448-D routers at the [edit interfaces *interface-name* optics-options] hierarchy level.

Table 75: Statements Supported on ACX6360 and ACX5448-D Routers

| Statement         | Options         | Release           | Interfaces Supported |
|-------------------|-----------------|-------------------|----------------------|
| fec               | sdsec   sdsec15 | 18.3R1, 19.2R1-S1 | ot                   |
| high-polarization | -               | 18.2R1, 19.2R1-S1 | ot                   |

Table 75: Statements Supported on ACX6360 and ACX5448-D Routers (*continued*)

| Statement                                                   | Options                                                        | Release           | Interfaces Supported |
|-------------------------------------------------------------|----------------------------------------------------------------|-------------------|----------------------|
| <code>laser-enable</code><br>  <code>no-laser-enable</code> | -                                                              | 18.2R1, 19.2R1-S1 | ot                   |
| <code>los-alarm-threshold</code>                            | -                                                              | 18.2R1, 19.2R1-S1 | ot                   |
| <code>los-warn-threshold</code>                             | -                                                              | 18.2R1, 19.2R1-S1 | ot                   |
| <code>modulation-format</code>                              | ( <code>16qam</code>   <code>8qam</code>   <code>qpsk</code> ) | 18.2R1, 19.2R1-S1 | ot                   |
| <code>signal-degrade</code>                                 | <code>ber-threshold-clear value</code>                         | 18.3R1, 19.2R1-S1 | ot                   |
|                                                             | <code>ber-threshold-signal-degrade value</code>                | 18.3R1, 19.2R1-S1 |                      |
|                                                             | <code>interval value</code>                                    | 18.3R1, 19.2R1-S1 |                      |
|                                                             | <code>q-threshold-signal-degrade</code>                        | 18.3R1, 19.2R1-S1 |                      |
|                                                             | <code>q-threshold-signal-degrade-clear</code>                  | 18.3R1, 19.2R1-S1 |                      |

Table 75: Statements Supported on ACX6360 and ACX5448-D Routers (*continued*)

| Statement | Options                                                                                            | Release           | Interfaces Supported |
|-----------|----------------------------------------------------------------------------------------------------|-------------------|----------------------|
| tca       | carrier-frequency-offset-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 | ot                   |
|           | carrier-frequency-offset-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 18.2R1, 19.2R1-S1 |                      |
|           | fec-ber<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)                           | 18.2R1, 19.2R1-S1 |                      |
|           | fec-corrected-errors-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.2R1, 19.2R1-S1 |                      |
|           | fec-uncorrected-words-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.2R1, 19.2R1-S1 |                      |
|           | laser-frequency-error-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.2R1, 19.2R1-S1 |                      |
|           | laser-frequency-error-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.2R1, 19.2R1-S1 |                      |
|           | pam-histogram-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)            | 18.2R1, 19.2R1-S1 |                      |
|           | residual-isi-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)             | 18.2R1, 19.2R1-S1 |                      |
|           | residual-isi-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)              | 18.2R1, 19.2R1-S1 |                      |
|           |                                                                                                    | 18.2R1, 19.2R1-S1 |                      |

Table 75: Statements Supported on ACX6360 and ACX5448-D Routers (*continued*)

| Statement  | Options                                                                               | Release           | Interfaces Supported |
|------------|---------------------------------------------------------------------------------------|-------------------|----------------------|
|            | rx-power-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    |                   |                      |
|            | rx-power-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.2R1, 19.2R1-S1 |                      |
|            | snr-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)          | 18.2R1, 19.2R1-S1 |                      |
|            | tec-current-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 |                      |
|            | tec-current-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 18.2R1, 19.2R1-S1 |                      |
|            | temperature-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.2R1, 19.2R1-S1 |                      |
|            | temperature-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 18.2R1, 19.2R1-S1 |                      |
|            | tx-power-high-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.2R1, 19.2R1-S1 |                      |
|            | tx-power-low-tca<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.2R1, 19.2R1-S1 |                      |
| tx-power   | dbm                                                                                   | 18.2R1, 19.2R1-S1 | ot                   |
| wavelength | nm                                                                                    | 18.2R1, 19.2R1-S1 | ot                   |



SEE ALSO

| [Configuring OTN](#) | 508**Supported Optics Options on PTX10008 and PTX10016 Series Routers**

[Table 76 on page 493](#) lists the statements that are supported on PTX10008 and PTX10016 Series routers at the `[edit interfaces interface-name optics-options]` hierarchy level.

**Table 76: Statements Supported on PTX10008 and PTX10016 Series Routers**

| Statement                                                           | Options                                                    | Release | Interfaces Supported |
|---------------------------------------------------------------------|------------------------------------------------------------|---------|----------------------|
| alarm low-light-alarm                                               | link-down   syslog                                         | 18.3R1  | ot                   |
| tca<br>carrier-frequency-offset-high-tca                            | (enable-tca   no-enable-tca   threshold   threshold-24hrs) | 18.3R1  | ot                   |
| tx-power                                                            | dbm                                                        | 18.3R1  | ot                   |
| warning<br>low-light-warning                                        | link-down   syslog                                         | 18.3R1  | ot                   |
| <a href="#">laser-enable</a><br>  <a href="#">no-laser-enable</a>   | -                                                          | 18.3R1  | ot                   |
| <a href="#">line-loopback</a><br>  <a href="#">no-line-loopback</a> | -                                                          | 18.3R1  | ot                   |
| <a href="#">prbs</a>   <a href="#">no-prbs</a>                      | -                                                          | 18.3R1  | ot                   |
| <a href="#">signal-degrade</a>                                      | <a href="#">ber-threshold-clear</a> <i>value</i>           | 18.3R1  | ot                   |
|                                                                     | <a href="#">ber-threshold-signal-degrade</a> <i>value</i>  | 18.3R1  |                      |
|                                                                     | <a href="#">interval</a> <i>value</i>                      | 18.3R1  |                      |

Table 76: Statements Supported on PTX10008 and PTX10016 Series Routers (*continued*)

| Statement | Options                                                                         | Release | Interfaces Supported |
|-----------|---------------------------------------------------------------------------------|---------|----------------------|
| tca       | odu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1  | ot                   |
|           | odu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1  |                      |
|           | odu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.3R1  |                      |
|           | odu-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  | 18.3R1  |                      |
|           | odu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | odu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1  |                      |
|           | odu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | otu-tca-bbe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | otu-tca-bbe-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1  |                      |
|           | otu-tca-es<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)     | 18.3R1  |                      |
|           |                                                                                 | 18.3R1  |                      |

Table 76: Statements Supported on PTX10008 and PTX10016 Series Routers (*continued*)

| Statement | Options                                                                         | Release | Interfaces Supported |
|-----------|---------------------------------------------------------------------------------|---------|----------------------|
|           | otu-tca-es-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)  |         |                      |
|           | otu-tca-ses<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | otu-tca-ses-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1  |                      |
|           | otu-tca-uas<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs)    | 18.3R1  |                      |
|           | otu-tca-uas-fe<br>(enable-tca   no-enable-tca  <br>threshold   threshold-24hrs) | 18.3R1  |                      |

Table 76: Statements Supported on PTX10008 and PTX10016 Series Routers (*continued*)

| Statement                               | Options                                                             | Release | Interfaces Supported |
|-----------------------------------------|---------------------------------------------------------------------|---------|----------------------|
| <code>trigger trigger-identifier</code> | <code>oc-lof</code><br>(hold-time (down   up)   ignore)             | 18.3R1  | ot                   |
|                                         | <code>oc-lom</code><br>(hold-time (down   up)   ignore)             | 18.3R1  |                      |
|                                         | <code>oc-los</code><br>(hold-time (down   up)   ignore)             | 18.3R1  |                      |
|                                         | <code>oc-tsf</code><br>(hold-time (down   up)   ignore)             | 18.3R1  |                      |
|                                         | <code>oc-wavelength-lock</code><br>(hold-time (down   up)   ignore) | 18.3R1  |                      |
|                                         | <code>odu-ais</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                      |
|                                         | <code>odu-bdi</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                      |
|                                         | <code>odu-bei</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                      |
|                                         | <code>odu-iae</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                      |
|                                         | <code>odu-lck</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                      |
|                                         | <code>odu-oci</code><br>(hold-time (down   up)   ignore)            | 18.3R1  |                      |
|                                         | <code>odu-sd</code> (hold-time<br>(down   up)   ignore)             | 18.3R1  |                      |
|                                         | <code>odu-ttim</code>                                               | 18.3R1  |                      |
|                                         | <code>opu-ptim</code><br>(hold-time (down   up)   ignore)           | 18.3R1  |                      |

Table 76: Statements Supported on PTX10008 and PTX10016 Series Routers (*continued*)

| Statement | Options                                                | Release | Interfaces Supported |
|-----------|--------------------------------------------------------|---------|----------------------|
|           | <b>otu-ais</b><br>(hold-time (down   up)   ignore)     | 18.3R1  |                      |
|           | <b>otu-bdi</b><br>(hold-time (down   up)   ignore)     | 18.3R1  |                      |
|           | <b>otu-fec-deg</b><br>(hold-time (down   up)   ignore) | 18.3R1  |                      |
|           | <b>otu-fec-exe</b><br>(hold-time (down   up)   ignore) | 18.3R1  |                      |
|           | <b>otu-iae</b><br>(hold-time (down   up)   ignore)     | 18.3R1  |                      |
|           | <b>otu-sd</b><br>(hold-time (down   up)   ignore)      | 18.3R1  |                      |
|           | <b>otu-ttim</b> (hold-time (down   up)   ignore)       | 18.3R1  |                      |

## RELATED DOCUMENTATION

| [Configuring OTN](#) | 508

## Forward Error Correction (FEC) and Bit Error Rate (BER)

### IN THIS SECTION

- [Understanding Pre-FEC BER Monitoring and BER Thresholds](#) | 498
- [Supported Forward Error Correction Modes on MX Series Routers](#) | 502
- [Supported Forward Error Correction Modes on PTX Series Routers](#) | 503
- [Supported Forward Error Correction Modes on ACX6360 Router](#) | 504

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” on page 504](#) 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 499](#). 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  $Q^2$ -factor measurements (also known as Q-factor).  $Q^2$ -factor is expressed in units of decibels relative to a  $Q^2$ -factor of zero (dBQ).  $Q^2$ -factor enables you to describe system margin in linear terms in contrast to BER values, which are nonlinear in nature. After you determine the thresholds, you must convert the threshold values from  $Q^2$ -factor to BER to enter them in the CLI by using scientific notation. BER can be converted to  $Q^2$ -factor by using the following equation:

$$Q^2\text{-factor} = 20 * \log_{10}(\sqrt{2} * \text{erfcinv}(2 * \text{BER}))$$

**TIP:** To convert between  $Q^2$ -factor and BER in a spreadsheet program, you can approximate the values by using the following formulas:

- To calculate  $Q^2$ -factor:

$$= 20 * \text{LOG10}(-\text{NORMSINV}(\text{BER}))$$

- To calculate BER:

$$= 1 - \text{NORMSDIST}(10^{(0.05 * Q^2\text{-factor})})$$

Table 77 on page 501 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 77: 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.975.11.4 (UFEC) | 9.1 dBQ                | 2.2E-3 | 10.1 dBQ                 | 6.9E-4 | 11.1 dBQ               | 1.6E-4 |
|                      | G.975.11.7 (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 78 on page 501 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 78: 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 |

Table 78: Example—Signal Degrade and Clear Thresholds After Configuration (*continued*)

| 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-24-10G-W-SFPP | G.975.11.4 (UFEC) | 9.1 dBQ                | 2.2E-3 | 12.1 dBQ                 | 2.8E-5 | 13.1 dBQ               | 3.1E-6 |
|                      | G.975.11.7 (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](#) | 424

## Supported Forward Error Correction Modes on MX Series Routers

[Table 79 on page 502](#) 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 79: FEC modes Supported on MX Series Routers

| Line Card    | FEC Mode                                                  | Port Speed | Junos Version |
|--------------|-----------------------------------------------------------|------------|---------------|
| MPC5E-40G10G | ( <i>gfec</i>   <i>efec</i>   <i>none</i>   <i>ufec</i> ) | 10G        | 13.3          |

Table 79: FEC modes Supported on MX Series Routers (*continued*)

| Line Card      | FEC Mode                                                  | Port Speed               | Junos Version |
|----------------|-----------------------------------------------------------|--------------------------|---------------|
| MPC5E-100G10G  | ( <i>gfec</i>   <i>efec</i>   <i>none</i>   <i>ufec</i> ) | 10G and 100G (GFEC only) | 13.3          |
| MIC6-10G-OTN   | ( <i>gfec</i>   <i>efec</i>   <i>none</i>   <i>ufec</i> ) | 10G                      | 13.3          |
| MIC6-100G-CFP2 | ( <i>gfec</i>   <i>none</i> )                             | 100G (GFEC only)         | 13.3          |
| MIC3-100G-DWDM | <i>gfec</i>   <i>hgfec</i>   <i>sdfec</i>                 | 100G                     | 15.1F5        |

SEE ALSO

| [fec](#) | [937](#)

## Supported Forward Error Correction Modes on PTX Series Routers

[Table 80 on page 503](#) lists the FEC modes that are supported on PTX Series routers at the `[edit interfaces interface-name otn-options]` hierarchy level.

Table 80: FEC Modes Supported on PTX Series Routers

| Line Card            | FEC Mode                                                  | Port Speed       | Junos Version                                     |
|----------------------|-----------------------------------------------------------|------------------|---------------------------------------------------|
| P1-PTX-24-10G-W-SFPP | ( <i>gfec</i>   <i>efec</i>   <i>none</i>   <i>ufec</i> ) | 10G              | 12.1X48, 12.3, 13.2 (PTX5000)<br>13.2R2 (PTX3000) |
| P2-10G-40G-QSFPP     | ( <i>gfec</i>   <i>efec</i>   <i>none</i>   <i>ufec</i> ) | 10G              | 14.1R2 (PTX5000)<br>15.1F6 (PTX3000)              |
| P2-100GE-OTN         | ( <i>gfec</i>   <i>none</i> )                             | 100G (GFEC only) | 14.1                                              |
| P1-PTX-2-100G-WDM    | ( <i>gfec-sdfec</i> )                                     | 100G             | 13.2 (PTX5000)<br>13.3 (PTX3000)                  |
| PTX-5-100G-WDM       | <i>gfec</i>   <i>sdfec</i>                                | 100G             | 15.1F6                                            |

SEE ALSO

Supported Forward Error Correction Modes on ACX6360 Router

[Table 81 on page 504](#) lists the FEC modes that are supported on ACX6360 routers at the [[edit interfaces interface-name optics-options](#)] hierarchy level.

Table 81: FEC modes Supported on ACX6360 Routers

| FEC Mode       | Modulation Format | Port Speed | Junos Version |
|----------------|-------------------|------------|---------------|
| <i>sdfec</i>   | QPSK              | 100G       | 18.3R1        |
| <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

Release History Table

| Release                | Description                                                                                                                                                                                                        |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">18.3R1</a> | 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 \(OTN\)](#) | [423](#)
- | [Configuring OTN](#) | [508](#)

## Dense Wavelength Division Multiplexing (DWDM) Interface Wavelength

### IN THIS SECTION

- [Ethernet DWDM Interface Wavelength Overview | 505](#)
- [Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength | 505](#)

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 | 984](#)

### 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 82 on page 506](#) shows configurable wavelengths and the corresponding frequency for each configurable wavelength.

**Table 82: 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          |
| 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          |

Table 82: Wavelength-to-Frequency Conversion Matrix (*continued*)

| Wavelength<br>(nm) | Frequency<br>(THz) | Wavelength<br>(nm) | Frequency<br>(THz) | Wavelength<br>(nm) | Frequency<br>(THz) |
|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
| 1534.64            | 195.35             | 1548.51            | 193.60             | 1562.64            | 191.85             |
| 1535.04            | 195.30             | 1548.91            | 193.55             | 1563.05            | 191.80             |
| 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             |
| 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](#) | [984](#)

## RELATED DOCUMENTATION

| [Understanding Optical Transport Network \(OTN\)](#) | [423](#)

| [Supported OTN and Optics Options](#) | [452](#)

## Configuring OTN

### IN THIS SECTION

- [Configuring OTN Interfaces on P1-PTX-24-10G-W-SFPP PIC](#) | [508](#)
- [Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | [512](#)
- [Configuring OTN Interfaces on P2-100GE-OTN PIC](#) | [516](#)
- [Configuring OTN Interfaces on MIC3-100G-DWDM MIC](#) | [521](#)
- [Configuring OTN Interfaces on PTX-5-100G-WDM PIC](#) | [526](#)
- [Configuring OTN Interface Options on PTX10K-LC1104](#) | [532](#)

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-tsfc | 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 | 426](#)

[optics-options | 1297](#)

[otn-options | 1299](#)

[signal-degrade | 968](#)

[preemptive-fast-reroute | 964](#)

## 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](#)” on [page 424](#).

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](#)” on page 498 for more information about pre-FEC BER monitoring and determining BER threshold settings.

## SEE ALSO

---

[optics-options | 1297](#)


---

[otn-options | 1299](#)


---

[signal-degrade | 968](#)


---

[preemptive-fast-reroute | 964](#)


---

**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-tsfc | 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 | 1297](#)

[otn-options | 1299](#)

[signal-degrade | 968](#)

[preemptive-fast-reroute | 964](#)

## 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)
```

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-tsfc | 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 (HGFE) 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 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 | 437](#)

[optics-options | 1297](#)

[otn-options | 1299](#)

[signal-degrade | 968](#)

[preemptive-fast-reroute | 964](#)

## 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 then 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-tsfc | 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 | 441](#)

*Upgrading Firmware on the 5-Port 100-Gigabit DWDM OTN PIC (PTX-5-100G-WDM)*

[optics-options | 1297](#)

[otn-options | 1299](#)

[signal-degrade | 968](#)

[preemptive-fast-reroute | 964](#)

## 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-tsfc | 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 (HGFE) 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 | 1297](#)

[otn-options | 1299](#)

[signal-degrade | 968](#)

[preemptive-fast-reroute | 964](#)

[Understanding the PTX10K-LC1104 Line Card | 445](#)

#### Release History Table

| Release                | Description                                                                                                                                        |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">15.1F6</a> | 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

|                                                                     |
|---------------------------------------------------------------------|
| <a href="#">Understanding Optical Transport Network (OTN)   423</a> |
| <a href="#">Supported OTN and Optics Options   452</a>              |

## ODU Path Delay Measurement for Performance Monitoring

IN THIS SECTION

- [Understanding ODU Path Delay Measurement | 537](#)
- [Enabling ODU Path Delay Measurement | 539](#)
- [Disabling ODU Path Delay Measurement | 541](#)

Use this topic to understand about

### Understanding ODU Path Delay Measurement

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 | 424](#)

[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 | 521](#)

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 526](#)

[remote-loop-enable | 966](#)

[Understanding the MIC3-100G-DWDM MIC | 437](#)

[Understanding the PTX-5-100G-WDM PIC | 441](#)

## 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 | 424](#)

[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 | 521](#)

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 526](#)  
[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 526](#)

[remote-loop-enable | 966](#)

[Understanding the MIC3-100G-DWDM MIC | 437](#)

[Understanding the PTX-5-100G-WDM PIC | 441](#)



## 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 | 424](#)

[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 | 521](#)

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 526](#)

[remote-loop-enable | 966](#)

[Understanding the MIC3-100G-DWDM MIC | 437](#)

[Understanding the PTX-5-100G-WDM PIC | 441](#)

## RELATED DOCUMENTATION

[Understanding Optical Transport Network \(OTN\) | 423](#)

[Supported OTN and Optics Options | 452](#)

# 3

PART

## Operation, Administration, and Management (OAM) for Ethernet Interfaces

---

Configuring Connectivity Fault Management | **545**

Configuring Link Fault Management | **655**

Configuring ITU-TY.1731 Ethernet Service OAM | **689**

---



# Configuring Connectivity Fault Management

## IN THIS CHAPTER

- [Ethernet Operations, Administration, and Maintenance | 545](#)
- [Introduction to OAM Connectivity Fault Management \(CFM\) | 546](#)
- [Configuring Connectivity Fault Management \(CFM\) | 553](#)
- [CFM Monitoring between CE and PE Devices | 624](#)
- [Configuring Continuity Check Messages | 648](#)

## Ethernet Operations, Administration, and Maintenance

This topic provides an overview to help you effectively configure Ethernet Operations, Administration, and Maintenance (OAM) on a network of Juniper Networks® MX Series 5G Universal Routing Platforms. For more information about configuring OAM parameters on Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

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

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

OAM can provide simple link-level information, provide performance statistics, or track end-to-end connectivity across the network. Simple link fault management (LFM) for Ethernet links is defined in IEEE 802.3ah.

IEEE 802.1ag OAM is supported on untagged, single tagged, and stacked VLAN interfaces.

Ethernet OAM functions are implemented as:

- Fault detection and notification (provided by continuity check messages)

- Path discovery (provided by the linktrace protocol)
- Fault isolation, verification, and recovery (isolation and verification are provided by a combination of protocols, while recovery is the function of protocols such as spanning tree)

The loopback protocol used in Ethernet OAM is modeled on the standard IP ping. After a fault is detected, the loopback protocol performs fault verification and isolation under the direction of a network operator.

The loopback is performed using request and response message pairs. A unicast loopback message is generated by a maintenance endpoint (MEP), and a loopback reply is generated by the destination maintenance intermediate point (MIP) or MEP.

The target MAC address is learned by the continuity check protocol or linktrace protocol. The loopback message's packet is always forwarded to a unique port by the originating MEP, as determined by a MAC table lookup or the MEP interface MAC address.

The target MIP or MEP generates a unicast loopback reply in response to the received loopback message. The loopback message follows the same path as a data packet, and intermediate bridges simply forward the packet to the destination MIP or MEP.

## RELATED DOCUMENTATION

[Ethernet OAM Connectivity Fault Management](#)

[Example: Configuring Ethernet CFM on Bridge Connections](#)

[Example: Configuring Ethernet CFM on Physical Interfaces](#)

## Introduction to OAM Connectivity Fault Management (CFM)

### IN THIS SECTION

- [Ethernet OAM Connectivity Fault Management | 547](#)
- [IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)

Use this topic to understand more about the connectivity fault management (CFM).

## Ethernet OAM Connectivity Fault Management

The most complete connectivity fault management (CFM) is defined in IEEE 802.1ag. This topic emphasizes the use of CFM in a Metro Ethernet environment.

The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol. Similar to IP traceroute, this protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.
- Fault isolation using the loopback protocol. Similar to IP ping, this protocol works with the continuity check protocol during troubleshooting.

CFM partitions the service network into various administrative domains. For example, operators, providers, and customers might be part of different administrative domains.

Each administrative domain is mapped into one maintenance domain providing enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible. Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where outermost domains are assigned a higher level than the innermost domains.

Customer end points have the highest maintenance domain level. In a CFM maintenance domain, each service instance is called a maintenance association. A *maintenance association* can be thought as a full mesh of maintenance endpoints (MEPs) having similar characteristics. MEPs are active CFM entities generating and responding to CFM protocol messages.

There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs only respond to CFM messages).

MEPs can be *up MEPs* or *down MEPs*. A link can connect a MEP at level 5 to a MEP at level 7. The interface at level 5 is an up MEP (because the other end of the link is at MEP level 7), and the interface at level 7 is a down MEP (because the other end of the link is at MEP level 5).

In a Metro Ethernet network, CFM is commonly used at two levels:

- By the service provider to check the connectivity among its provider edge (PE) routers
- By the customer to check the connectivity among its customer edge (CE) routers

**NOTE:** The configured customer CFM level must be greater than service provider CFM level.

In many Metro Ethernet networks, CFM is used to monitor connectivity over a VPLS and bridge network.

**NOTE:** In ACX Series routers, OAM for VPLS is supported only on ACX5048, ACX5096, and ACX5448 routers.

#### SEE ALSO

[CFM Monitoring between CE and PE Devices | 624](#)

[Configuring Continuity Check Messages | 648](#)

## IEEE 802.1ag OAM Connectivity Fault Management Overview

### IN THIS SECTION

- [OAM Connectivity Fault Management Overview | 548](#)
- [Connectivity Fault Management Key Elements | 550](#)
- [Best Practices for Configuring 802.1ag Ethernet OAM for VPLS | 552](#)

Ethernet interfaces on M7i and M10i routers with the Enhanced CFEB (CFEB-E) and on M120, M320, MX Series, T Series, and PTX Series routers support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM). The goal of CFM is to monitor an Ethernet network that may comprise one or more service instances. Junos OS supports IEEE 802.1ag connectivity fault management.

**NOTE:** MX Series Virtual Chassis does not support distributed inline connectivity fault management.

ACX Series routers support CFM on aggregated Ethernet interfaces with continuity check interval of 100 milliseconds or higher.

### ***OAM Connectivity Fault Management Overview***

In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces. Connectivity fault management (CFM) sessions operate in distributed mode and are processed on the Flexible PIC Concentrator



(FPC) on aggregated Ethernet interfaces. As a result, graceful Routing Engine switchover (GRES) is supported on aggregated Ethernet interfaces. In releases before Junos OS Release 13.3, CFM sessions operate in centralized mode and are processed on the Routing Engine. However, CFM sessions are not supported on aggregated Ethernet interfaces if the interfaces that form the aggregated Ethernet bundle are in mixed mode. CFM sessions with a continuity check message (CCM) interval of 10 milliseconds are not supported over aggregated Ethernet interfaces.

CFM sessions are distributed by default. All CFM sessions must operate in either only distributed or only centralized mode. A mixed operation of distributed and centralized modes for CFM sessions is not supported. To disable the distribution of CFM sessions on aggregated Ethernet interfaces and make the sessions operate in centralized mode, include the **no-aggregate-delegate-processing** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

**NOTE:** As a requirement for Ethernet OAM 802.1ag to work, distributed periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine. You can only disable PPM on the Packet Forwarding Engine. To disable PPM on the PFE, include the **ppm no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

**NOTE:**

- CFM sessions are supported on aggregated Ethernet interfaces if the interfaces that form the aggregated Ethernet bundle are in mixed mode when the **no-aggregate-delegate-processing** command is enabled.
- Starting in Junos OS Release 14.2, for CFM sessions in centralized mode, we recommend that you configure a maximum of 40 CFM sessions with continuity check message (CCM) interval of 100 milliseconds (100 ms) or a maximum of 400 CFM sessions with CCM interval of 1 second (1 s). If CFM sessions are configured beyond this limit, CFM might not work as expected. You might observe issues when the state of multiple links change or when the line cards are restarted.

Note that these limits have been derived by considering a protocol data unit (PDU) load of 400 packets per second (pps) on the Routing Engine. This limit varies depending on the Routing Engine load. If the Routing Engine experiences heavy load, expect some variations to this limit.

Starting in Junos OS Release 10.3, on interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces. Starting in Junos OS Release 12.3, CFM does not support Multichassis Link Aggregation (MC-LAG). Do not configure the **mc-ae** statement when you configure CFM.

Starting in Junos OS Release 11.3, on T Series and M320 routers, CFM is not supported on interfaces configured with CCC encapsulation. If you configure CFM, the system displays the following message: **“MEPs cannot be configured on ccc interface on this platform”**.

Network entities such as operators, providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management, perform end-to-end monitoring, and still avoid security breaches.

Starting in Junos OS Release 17.4, you can enable support for IEEE 802.1ag CFM on pseudowire service interfaces by configuring maintenance intermediate points (MIPs) on the pseudowire service interfaces. Pseudowire service interfaces support configuring of subscriber interfaces over MPLS pseudowire termination. Termination of subscriber interfaces over PW enables network operators to extend their MPLS domain from the Access/Aggregation network to the service edge and use uniform MPLS label provisioning for a larger portion of their network.

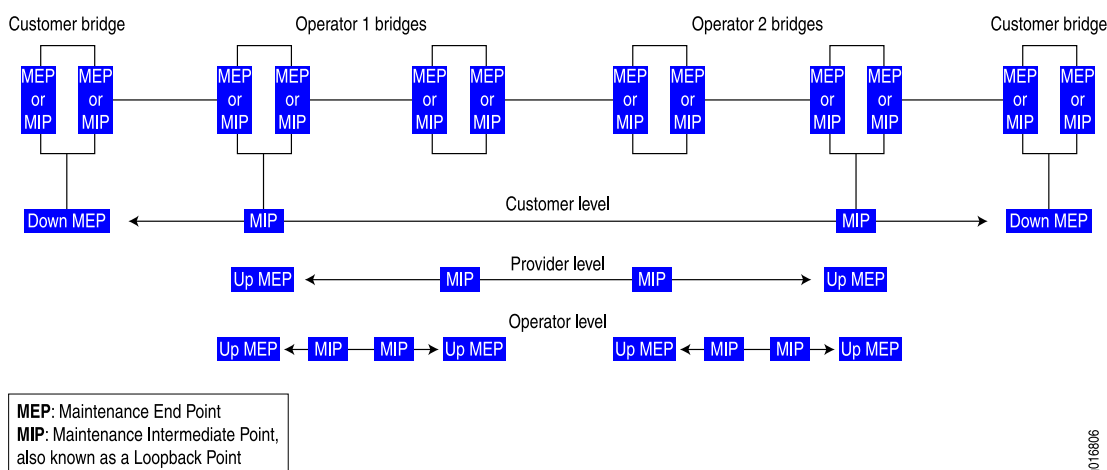
**NOTE:** The CFM MIP session is supported only on the pseudowire services interface and not on the pseudowire services tunnel interface.

IEEE 802.1ag OAM supports graceful Routing Engine switchover (GRES). IEEE 802.1ag OAM is supported on untagged, single tagged, and stacked VLAN interfaces.

### Connectivity Fault Management Key Elements

Figure 8 on page 550 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

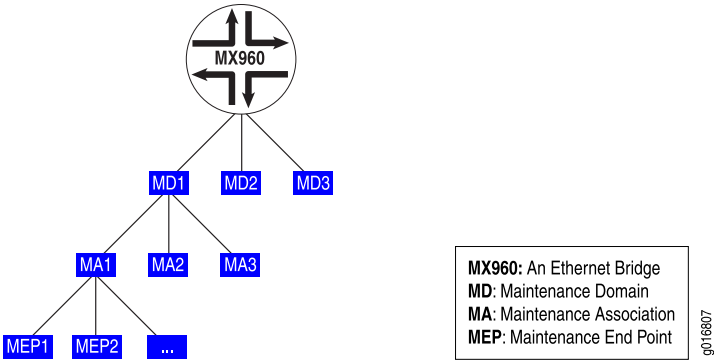
**Figure 8: Relationship Among MEPs, MIPs, and Maintenance Domain Levels**



**NOTE:** On ACX Series routers, the maintenance intermediate points (MIP) are supported only on the ACX5048 and ACX5096 routers.

A maintenance association is a set of MEPs configured with the same maintenance association identifier and maintenance domain level. [Figure 9 on page 551](#) shows the hierarchical relationships between the Ethernet bridge, maintenance domains, maintenance associations, and MEPs.

**Figure 9: Relationship Among Bridges, Maintenance Domains, Maintenance Associations, and MEPs**



## Best Practices for Configuring 802.1ag Ethernet OAM for VPLS

**BEST PRACTICE:** The logical interfaces in a VPLS routing instance may have the same or different VLAN configurations. VLAN normalization is required to switch packets correctly among these interfaces. VLAN normalization is effectively VLAN translation wherein the VLAN tags of the received packet need to be translated if they are different than the normalized VLAN tags.

For MX Series routers, the normalized VLAN is specified using one of the following configuration statements in the VPLS routing instance:

- **vlan-id *vlan-number***
- **vlan-id none**
- **vlan-tags outer *outer-vlan-number* inner *inner-vlan-number***

You must configure **vlan-maps** explicitly on all interfaces belonging to the routing instance.

The following forwarding path considerations must be observed:

- Packet receive path:
  - This is the forwarding path for packets received on the interfaces.
  - 802.1ag Ethernet OAM for VPLS uses implicit interface filters and forwarding table filters to flood, accept, and drop the CFM packets.
- Packet transmit path:
  - The JUNOS Software uses the router's hardware-based forwarding for CPU-generated packets.
  - For Down MEPs, the packets are transmitted on the interface on which the MEP is configured.
  - In MX series routers, for Up MEPs, the packet must be flooded to other interfaces in the VPLS routing instance. The router creates a flood route tied to a flood next hop (with all interfaces to flood) and then sources the packet to be forwarded with this flood route.
  - The router also uses implicit-based forwarding for CPU generated packets. The result is for the flood next hop tied to the flood route to be tied to the filter term. The filter term uses match criteria to correctly identify the host-generated packets.

## SEE ALSO

[connectivity-fault-management](#) | 1132

[Configuring Port Status TLV and Interface Status TLV](#) | 626

**Release History Table**

| Release                | Description                                                                                                                                                                                                                                                                                                              |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">17.4R1</a> | Starting in Junos OS Release 17.4, you can enable support for IEEE 802.1ag CFM on pseudowire service interfaces by configuring maintenance intermediate points (MIPs) on the pseudowire service interfaces.                                                                                                              |
| <a href="#">14.2</a>   | Starting in Junos OS Release 14.2, for CFM sessions in centralized mode, we recommend that you configure a maximum of 40 CFM sessions with continuity check message (CCM) interval of 100 milliseconds (100 ms) or a maximum of 400 CFM sessions with CCM interval of 1 second (1 s).                                    |
| <a href="#">12.3</a>   | Starting in Junos OS Release 12.3, CFM does not support Multichassis Link Aggregation (MC-LAG). Do not configure the <b>mc-ae</b> statement when you configure CFM.                                                                                                                                                      |
| <a href="#">11.3</a>   | Starting in Junos OS Release 11.3, on T Series and M320 routers, CFM is not supported on interfaces configured with CCC encapsulation.                                                                                                                                                                                   |
| <a href="#">10.3</a>   | Starting in Junos OS Release 10.3, on interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces. |
| <a href="#">9.3</a>    | In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces.                                                                                                                                                                                                                                     |

**RELATED DOCUMENTATION**

[CFM Monitoring between CE and PE Devices | 624](#)

[Configuring Continuity Check Messages | 648](#)

**Configuring Connectivity Fault Management (CFM)****IN THIS SECTION**

- [Creating a Maintenance Domain | 555](#)
- [Configuring Maintenance Intermediate Points \(MIPs\) | 556](#)

- [Configuring Maintenance Association Intermediate Points in ACX Series | 557](#)
- [Creating a Maintenance Association | 561](#)
- [Continuity Check Protocol Parameters Overview | 562](#)
- [Configuring Continuity Check Protocol Parameters for Fault Detection | 563](#)
- [Configuring a MEP to Generate and Respond to CFM Protocol Messages | 564](#)
- [Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements | 568](#)
- [Configuring Service Protection for VPWS over MPLS Using the MEP Interface | 570](#)
- [Configuring Linktrace Protocol in CFM | 576](#)
- [Configuring Rate Limiting of Ethernet OAM Messages | 576](#)
- [Configuring Ethernet Local Management Interface | 579](#)
- [Configuring a CFM Action Profile to Specify CFM Actions for CFM Events | 589](#)
- [CFM Action Profile to Bring Down a Group of Logical Interfaces Overview | 591](#)
- [Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces | 592](#)
- [Enabling Enhanced Connectivity Fault Management Mode | 596](#)
- [Configuring M120 and MX Series Routers for CCC Encapsulated Packets | 597](#)
- [Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades | 599](#)
- [Configuring Unified ISSU for 802.1ag CFM | 600](#)
- [Junos OS Support for Performance Monitoring Compliant with Technical Specification MEF 36 | 604](#)
- [Damping CFM performance Monitoring Traps and Notifications to Prevent Congestion of The NMS | 605](#)
- [Example: Configuring Ethernet CFM on Physical Interfaces | 606](#)
- [Example: Configuring Ethernet CFM on Bridge Connections | 609](#)
- [Example: Configuring Ethernet CFM over VPLS | 614](#)

Use this topic to configure connectivity fault management features such as maintenance domains, maintenance associations, maintenance intermediate points (MIPs), and continuity check parameters. You can also use this topic to configure an action profile to specify the CFM action that must be performed when a specific CFM event occurs.

## Creating a Maintenance Domain

To enable connectivity fault management (CFM) on an Ethernet interface, you must first configure a maintenance domain and specify the name of the maintenance domain. You can also specify the format of the name. For instance, if you specify the name format to be domain name service (DNS) format, you can specify the name of the maintenance domain as `www.juniper.net`. The default name format is ASCII character string.

**NOTE:** For logical interfaces, the maintenance domain name must be unique across logical systems. If you configure the same maintenance domain name across logical systems, then you receive the following error message: **error: configuration check-out failed**.

During the creation of the maintenance domain, you can also specify the maintenance domain level. The maintenance domain level indicates the nesting relationship between various maintenance domains. The maintenance domain level is embedded in each of the CFM frames.

To create a maintenance domain:

1. In configuration mode, create a maintenance domain by specifying the name and the name format at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set maintenance-domain md-name name-format option
```

**NOTE:** If you configure the maintenance domain name length greater than 45 octet, then the following error message is displayed: **error: configuration check-out failed**.

2. Specify the maintenance domain level by specifying the value at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set maintenace-domain md-name level number
```

## SEE ALSO

[connectivity-fault-management](#) | 1132

[maintenance-domain](#) | 1051

[name-format](#) | 1058

[level](#) | 1043

## Configuring Maintenance Intermediate Points (MIPs)

MX Series routers support maintenance intermediate points (MIPs) for the Ethernet OAM 802.1ag CFM protocol at a bridge-domain level. This enables you to define a maintenance domain for each default level. The MIPs names are created as **default-level-number** at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain]** hierarchy level. Use the **bridge-domain**, **instance**, **virtual-switch**, and **mip-half-function** MIP options to specify the MIP configuration.

Use the **show oam ethernet connectivity-fault-management mip (bridge-domain | instance-name | interface-name)** command to display the MIP configurations.

To configure the maintenance intermediate point (MIP):

1. Configure a bridge domain under a user-defined virtual switch by specifying the **virtual-switch** statement and the name of the user-defined virtual switch, at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name default-x]** hierarchy level.

**NOTE:** A bridge domain must be specified by name only if it is configured by including the **vlan-id** statement under the **virtual-switch** statement. If a bridge domain is configured with a range of VLAN IDs, then the VLAN IDs must be explicitly listed after the bridge domain name.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name default-x]
user@host# set virtual-switch virtual-switch-name bridge-domain bridge-domain-name vlan-id value
```

**NOTE:** You can also configure the bridge domain for the default virtual switch by including the **bridge-domain** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]** hierarchy level.

2. Configure the VPLS routing instance for the default maintenance domain.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
user@host# set instance instance-name
```



3. Configure the maintenance intermediate point (MIP) half function to divide the MIP functionality into two unidirectional segments to improve network coverage by increasing the number of MIPs that are monitored. The MIP half function also responds to loop-back and link-trace messages to identify faults.

**NOTE:** Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations be the same.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name default- x]
user@host# set mip-half-function (none | default | explicit)
```

#### SEE ALSO

[bridge-domain | 1125](#)

[connectivity-fault-management | 1132](#)

[instance | 1032](#)

[mip-half-function | 1057](#)

[virtual-switch | 1423](#)

## Configuring Maintenance Association Intermediate Points in ACX Series

### IN THIS SECTION

- [Configuring the Maintenance Domain Bridge Domain | 558](#)
- [Configuring the Maintenance Domain MIP Half Function | 559](#)
- [Configuring the Maintenance Association Intermediate Points with Bridge Domain | 559](#)
- [Configuring the Maintenance Association Intermediate Points with Circuit Cross-Connect | 559](#)
- [Configuring the Maintenance Association Intermediate Points with Bridge Domain when Maintenance Association End Point is Configured | 560](#)
- [Configuring the Maintenance Association Intermediate Points with Circuit Cross-Connect when Maintenance Association End Point is Configured | 560](#)

Maintenance Intermediate Point (MIP) provides monitoring capability of intermediate points for services such as Layer 2 bridging, Layer 2 circuit, and Layer 2 VPN. ACX5048 and ACX5096 routers support MIPs for the Ethernet OAM 802.1ag CFM protocol. Use the bridge-domain, interface, and mip-half-function MIP options to specify the MIP configuration.

**NOTE:** ACX5048 and ACX5096 routers do not support MIP configuration on VPLS services.

**NOTE:** ACX5448 router do not support MIP.

**NOTE:** Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations be the same.

To display MIP configurations, use the **show oam ethernet connectivity-fault-management mip (bridge-domain | instance-name | interface-name)** command.

The following MIP configurations are supported in ACX5048 and ACX5096 routers:

- MIP with with bridge domain
- MIP with circuit cross-connect (CCC)
- MIP with bridge domain when maintenance association end point is configured
- MIP with CCC when maintenance association end point is configured

The following sections describe MIP configuration:

#### ***Configuring the Maintenance Domain Bridge Domain***

To configure the bridge domain, include the **vlan** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain maintenance-domain-name]** hierarchy level.

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

### **Configuring the Maintenance Domain MIP Half Function**

MIP Half Function (MHF) divides MIP functionality into two unidirectional segments, improves visibility with minimal configuration, and improves network coverage by increasing the number of points that can be monitored. MHF extends monitoring capability by responding to loopback and linktrace messages to help isolate faults.

Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the *MIP half function* value for all maintenance domains and maintenance associations be the same. To configure the MIP half function, include the [mip-half-function](#) statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain maintenance-domain-name]** hierarchy level.

### **Configuring the Maintenance Association Intermediate Points with Bridge Domain**

In ACX5048 and ACX5096 routers, you can configure the MIP with bridge domain. The following is a sample to configure the MIP with bridge domain:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain default-6 {
        vlan bd1;
        mip-half-function default;
      }
    }
  }
}
```

### **Configuring the Maintenance Association Intermediate Points with Circuit Cross-Connect**

In ACX5048 and ACX5096 routers, you can configure the MIP with circuit cross-connect (CCC). The following is a sample to configure the MIP with CCC:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain default-6 {
        interface xe-0/0/42.0;
        mip-half-function default;
      }
    }
  }
}
```

### ***Configuring the Maintenance Association Intermediate Points with Bridge Domain when Maintenance Association End Point is Configured***

In ACX5048 and ACX5096 routers, you can configure the MIP with bridge domain when a maintenance association end point (MEP) is configured. The following is a sample to configure the MIP with bridge domain when MEP is configured:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md2 {
        level 5;
        mip-half-function default;
        maintenance-association ma2 {
          continuity-check {
            interval 1s;
          }
          mep 222 {
            interface xe-0/0/42.0;
            direction up;
          }
        }
      }
    }
  }
}
```

### ***Configuring the Maintenance Intermediate Points with Circuit Cross-Connect when Maintenance Association End Point is Configured***

In ACX5048 and ACX5096 routers, you can configure the MIP with circuit cross-connect (CCC) when a maintenance association end point (MEP) is configured. The following is a sample to configure the MIP with CCC when MEP is configured:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md2 {
        level 5;
        mip-half-function default;
        maintenance-association ma2 {
          continuity-check {
            interval 1s;
          }
        }
      }
    }
  }
}
```

```

        mep 222 {
            interface xe-0/0/42.0;
            direction up;
        }
    }
}
}
}
}

```

## SEE ALSO

[bridge-domain | 1125](#)

[connectivity-fault-management | 1132](#)

[instance | 1032](#)

[mip-half-function | 1057](#)

## Creating a Maintenance Association

To create a maintenance association, include the **maintenance-association *ma-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*]** hierarchy level.

Maintenance association names can be in one of the following formats:

- As a plain ASCII character string
- As the VLAN identifier of the VLAN you primarily associate with the maintenance association
- As a two-octet identifier in the range from 0 through 65,535
- As a name in the format specified by RFC 2685

The default short name format is an ASCII character string.

To configure the maintenance association short name format, include the **short-name-format (*character-string* | *vlan* | *2octet* | *rfc-2685-vpn-id*)** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level.

## SEE ALSO

[connectivity-fault-management](#) | [1132](#)

[name-format](#) | [1058](#)

[short-name-format](#) | [1076](#)

## Continuity Check Protocol Parameters Overview

The continuity check protocol is used for fault detection by maintenance end points (MEPs) within a maintenance association. The MEP periodically sends continuity check multicast messages. The continuity check protocol packets use the ethertype value 0x8902 and the multicast destination MAC address 01:80:c2:00:00:32.

The following list describes the continuity check protocol parameters you can configure:

- **interval**—Frequency of the continuity check messages (CCM) i.e time between the transmission of the CCM messages. You can specify 10 minutes (**10m**), 1 minute (**1m**), 10 seconds (**10s**), 1 second (**1s**), 100 milliseconds (**100ms**), or 10 milliseconds (**10ms**). The default value is 1 minute. For instance, if you specify the interval as 1 minute, the MEP sends the continuity check messages every minute to the receiving MEP.

**NOTE:** For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine by default. You can only disable PPM on the Packet Forwarding Engine. To disable PPM on the Packet Forwarding Engine, use the **no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

Continuity check interval of 10 milliseconds is not supported for CFM sessions over a label-switched interface (LSI).

- **hold-interval**—Frequency at which the MEP database can be flushed, if no updates occur. Receiving MEPs use the continuity check messages to build a MEP database of all MEPs in the maintenance association. The frequency is the number of minutes to wait before flushing the MEP database if no updates occur. The default value is 10 minutes.

**NOTE:** Hold timer based flushing is applicable only for autodiscovered remote MEPs and not for statically configured remote MEPs.

The hold interval logic runs a polling timer per CFM session level (not per remote MEP level) where the polling timer duration is equal to the configured hold time. When the polling timer expires, it deletes all the autodiscovered remote MEP entries which have been in the failed state for a time period equal to or greater than the configured hold time. If the remote MEP completes the hold time duration in the

failed state, then flushing will not occur until the next polling timer expires. Hence remote MEP flushing may not happen exactly at the configured hold time.

- **loss-threshold**—Number of continuity check messages that can be lost before the router marks the MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

#### SEE ALSO

[hold-interval](#) | 1031

[interval](#) | 1040

[loss-threshold](#) | 1047

### Configuring Continuity Check Protocol Parameters for Fault Detection

The continuity check protocol is used for fault detection by a maintenance association end point (MEP) within a maintenance association. A MEP periodically generates and responds to continuity check multicast messages. The continuity check protocol packets use the ethertype value 0x8902 and the multicast destination MAC address 01:80:c2:00:00:32. The receiving MEPs use the continuity check messages (CCMs) to build a MEP database of all MEPs in the maintenance association.

To configure continuity check protocol parameters:

1. Specify the time to wait in minutes before flushing the MEP database, if no updates occur, with a value from 1 minute through 30,240 minutes. The default value is 10 minutes.

**NOTE:** Flushing based on the hold timer is applicable only for autodiscovered remote MEPs and not for statically configured remote MEPs.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name continuity-check]
user@host# set hold-interval minutes
```

2. Specify the time to wait (duration) between the transmissions of CCMs. The duration can be one of the following values: 10 minutes (10m), 1 minute (1m), 10 seconds (10s), 1 second (1s), 100 milliseconds (100ms), or 10 milliseconds (10ms). The default value is 1 minute.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name continuity-check]
```

```
user@host# set interval duration
```

3. Specify the number of continuity check messages that can be lost before the router marks the MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
 maintenance-association ma-name continuity-check]
user@host# set loss-threshold number
```

## SEE ALSO

[continuity-check | 1011](#)

[hold-interval | 1031](#)

[interval | 1040](#)

[loss-threshold | 1047](#)

## Configuring a MEP to Generate and Respond to CFM Protocol Messages

### IN THIS SECTION

- [Configuring a Maintenance Association End Point \(MEP\) | 565](#)
- [Configuring a remote Maintenance Association End Point \(MEP\) | 567](#)

A maintenance association end point (MEP) refers to the boundary of a domain. A MEP generates and responds to connectivity fault management (CFM) protocol messages. You can configure multiple up MEPs for a single combination of maintenance association ID and maintenance domain ID for interfaces belonging to a particular VPLS service or a bridge domain. You can configure multiple down MEPs for a single instance of maintenance domain identifier and maintenance association name to monitor services provided by Virtual Private LAN service (VPLS), bridge domain, circuit cross-connect (CCC), or IPv4 domains.

For layer 2 VPNs routing instances (local switching) and EVPN routing instances, you can also configure multiple up MEPs for a single combination of maintenance association ID and maintenance domain ID on logical interfaces.. The logical interface can be configured on different devices or on the same device. To support multiple up MEPs on two IFLs, enhanced IP network services must be configured for the chassis.



You can enable automatic discovery of a MEP. With automatic discovery a MEP is enabled to accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association. If automatic discovery is not enabled, the remote MEPs must be configured. If the remote MEP is not configured, the CCMs from the remote MEP are treated as errors.

Continuity measurement is provided by an existing continuity check protocol. The continuity for every remote MEP is measured as the percentage of time that remote MEP was operationally up over the total administratively enabled time. Here, the operational uptime is the total time during which the CCM adjacency is active for a particular remote MEP and the administrative enabled time is the total time during which the local MEP is active. You can also restart the continuity measurement by clearing the currently measured operational uptime and the administrative enabled time.

### **Configuring a Maintenance Association End Point (MEP)**

To configure a maintenance association end point:

1. Specify an ID for the MEP at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]**. You can specify any value from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name]
user@host# set mep mep-id
```

2. Enable maintenance end point automatic discovery so the MEP can accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name mep mep-id]
user@host# set auto-discovery
```

3. Specify the direction in which the CCM packets are transmitted for the MEP. You can specify up or down. If you specify the direction as up, CCMs are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on the MEP. If you specify the direction as down, CCMs are transmitted only out of the interface configured on the MEP.

**NOTE:** Ports in the Spanning Tree Protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in an STP blocking state without the continuity check protocol configured do block CFM packets.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name mep mep-id]
user@host# set direction down
```

**NOTE:** Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 5G Universal Routing Platforms, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. For all other interfaces on MX Series routers and on all other routers and switches, you must continue to configure the **no-control-word** statement at the **[edit routing-instances *routing-instance-name* protocols l2vpn]** or **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** hierarchy level when you configure CFM MEPs. Otherwise, the CFM packets are not transmitted, and the **show oam ethernet connectivity-fault-management mep-database** command does not display any remote MEPs.

4. Specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface. On MX Series routers, the MEP can be attached to a specific VLAN of a trunk interface.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name mep mep-id]
user@host# set interface interface-name
```

5. Specify the IEEE 802.1 priority bits that are used by continuity check and link trace messages. You can specify a value from through 7 as the priority.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name mep mep-id]
user@host# set priority number
```

6. Specify the lowest priority defect that generates a fault alarm whenever CFM detects a defect. Possible values include: all -defects, err-xcon, mac-rem-err-xcon, no-defect, rem-err-xcon, and xcon.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name mep mep-id]
user@host# set lowest-priority-defect mac-rem-err-xcon
```

7. Specify the ID of the remote MEP at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]**. You can specify any value from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-namemep mep-id]
user@host# set remote-mep mep-id
```

## SEE ALSO

| [priority](#) | [1068](#)

### Configuring a remote Maintenance Association End Point (MEP)

To configure a remote maintenance association end point:

1. Configure the remote MEP by specifying the MEP ID at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*]**. You can specify any value from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-namemep mep-id]
user@host# edit remote-mep mep-id
```

2. Specify the name of the action profile to be used for the remote MEP by including the **action-profile *profile-name*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id* remote-mep *remote-mep-id*]**. The profile must be defined at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-namemep mep-id remote-mep remote-mep-id]
user@host# set action-profile profile-name
```

3. Configure the remote MEP to detect initial loss of connectivity. By default, the MEP does not generate loss-of-continuity (LOC) defect messages. When you configure the **detect-loc** statement, a loss-of-continuity (LOC) defect is detected if no continuity check message is received from the remote MEP within a period equal to 3.5 times the continuity check interval configured for the maintenance association. If a LOC defect is detected, a syslog error message is generated.

**NOTE:** When you configure connectivity-fault management (CFM) along with **detect-loc**, any **action-profile** configured to bring down the interface is executed if continuity check message is not received . However, the **action-profile** is not executed if you have not configured **detect-loc** and continuity check message is not received.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-namemep mep-id remote-mep remote-mep-id]
user@host# set detect-loc
```

SEE ALSO

[remote-mep | 1074](#)

SEE ALSO

[action-profile | 996](#)

[auto-discovery | 999](#)

[connectivity-fault-management | 1132](#)

[detect-loc | 1020](#)

[direction | 1022](#)

[lowest-priority-defect | 1048](#)

## Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging Service Level Agreements (SLAs). By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

Before you can perform Ethernet frame delay measurements on MX Series routers, you must have done the following:

- Configured Ethernet OAM and CFM correctly
- Prepared the measurement between two compatibly configured MX Series routers

- Enabled the distributed periodic packet management daemon (ppmd)
- Avoided trying to perform Ethernet frame delay measurement on aggregated Ethernet or pseudowire interfaces, which are not supported
- Made sure the hardware-assisted timestamping is supported if that feature is configured

At the end of this configuration, you create two MX Series routers that can perform and display Ethernet frame delay measurements on Ethernet interfaces using optional hardware timestamping. By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can optionally use hardware timing to assist in this process and increase the accuracy of the delay measurement results. This assistance is available on the reception path.

To configure hardware-assisted timestamping:

1. To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        performance-monitoring {
          hardware-assisted-timestamping; # Enable timestamping in hardware.
        }
      }
    }
  }
}
```

2. Ethernet frame delay measurement requires that distributed PPMD is enabled. Before you can gather statistics for Ethernet frame delay measurement, you must make sure that PPMD is configured properly. Without distributed PPMD, delay measurement results are not valid.

To perform Ethernet frame delay measurement, make sure that the following configuration statement is **NOT** present:

```
[edit routing-options]
ppm {
  no-delegate-processing; # This turns distributed PPMD OFF.
}
```

SEE ALSO

|                                                                                                    |
|----------------------------------------------------------------------------------------------------|
| Ethernet Frame Delay Measurements Overview   690                                                   |
| Triggering an Ethernet Frame Delay Measurements Session   723                                      |
| Viewing Ethernet Frame Delay Measurements Statistics   741                                         |
| Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces   728 |
| Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces   734 |
| Configuring ETH-DM with Untagged Interfaces                                                        |

### Configuring Service Protection for VPWS over MPLS Using the MEP Interface

You can enable service protection for a virtual private wire service (VPWS) over MPLS by specifying a working path or protect path on the MEP. Service protection provides end-to-end connection protection of the working path in the event of a failure.

To configure service protection, you must create two separate transport paths—a working path and a protect path. You can specify the working path and protect path by creating two maintenance associations. To associate the maintenance association with a path, you must configure the **interface** statement for the MEP within the maintenance association and specify the path as working or protect.

**NOTE:** If the path is not specified, the session monitors the active path.

Table 83 on page 570 describes the available service protection options.

Table 83: Service Protection Options

| Option         | Description                 |
|----------------|-----------------------------|
| <b>working</b> | Specifies the working path. |
| <b>protect</b> | Specifies the protect path. |

In this configuration, we enable service protection for the VPWS service. The CCM session is configured for the working path and references the CCM session configured for the protect path using the **protect-maintenance-association** statement. The name of the protect transport path for the maintenance association is configured and associated with the maintenance association for the working path.

To configure service protection for VPWS over MPLS:

1. In configuration mode, create a maintenance domain by specifying the name and the name format at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set maintenance-domain md-name name-format option
```

**NOTE:** If you configure the maintenance domain name length greater than 45 octet, then the following error message is displayed: **error: configuration check-out failed.**

2. Specify the maintenance domain level by specifying the value at the **[edit protocols oam ethernet connectivity-fault-management ]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set maintenace-domain md-name level number
```

3. Create a maintenance association for the working path by specifying the name and the short name format at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name*]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name]
user@host# set maintenance-association test-ma short-name-format option
```

4. Specify the maintenance association name used for connection protection and the name of the automatic-protection-switching profile (aps-profile) at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name]
user@host# set protect-maintenance-association ma-name aps-profile aps-profile-name
```

5. Specify the time to wait between transmissions of continuity check messages at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* continuity-check ]** hierarchy level. The duration can be one of the following values: 10 minutes(10m), 1 minute(1m), 10 seconds(10s), 1 second(1s), 100 milliseconds(100ms), or 10 milliseconds(10ms). The default value is 1 minute.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name continuity-check]
user@host# set interval option
```

6. Specify an ID for the MEP at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]**. You can specify any value from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name]
user@host# set mep mep-id
```

7. Enable maintenance end point automatic discovery so the MEP can accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
user@host# set auto-discovery
```

8. Specify the direction in which the CCM packets are transmitted for the MEP. You can specify up or down. If you specify the direction as up, CCMs are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on the MEP. If you specify the direction as down, CCMs are transmitted only out of the interface configured on the MEP.

**NOTE:** Ports in the Spanning Tree Protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in an STP blocking state without the continuity check protocol configured do block CFM packets.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
user@host# set direction down
```



**NOTE:** Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 5G Universal Routing Platforms, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. For all other interfaces on MX Series routers and on all other routers and switches, you must continue to configure the **no-control-word** statement at the **[edit routing-instances *routing-instance-name* protocols l2vpn]** or **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** hierarchy level when you configure CFM MEPs. Otherwise, the CFM packets are not transmitted, and the **show oam ethernet connectivity-fault-management mep-database** command does not display any remote MEPs.

9. Specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface. On MX Series routers, the MEP can be attached to a specific VLAN of a trunk interface. Also, specify the transport path as working.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
user@host# set interface interface-name working
```

10. Create a maintenance association for the protection path by specifying the name and the short name format at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name*]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name]
user@host# set maintenance-association ma-name short-name-format option
```

11. Specify the time to wait between transmissions of continuity check messages at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* continuity-check ]** hierarchy level. The duration can be one of the following values: 10 minutes(10m), 1 minute(1m), 10 seconds(10s), 1 second(1s), 100 milliseconds(100ms), or 10 milliseconds(10ms). The default value is 1 minute.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name continuity-check]
user@host# set interval option
```

12. Specify an ID for the MEP at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]**. You can specify any value from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name]
user@host# set mep mep-id
```

13. Enable maintenance end point automatic discovery so the MEP can accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
user@host# set auto-discovery
```

14. Specify the direction in which the CCM packets are transmitted for the MEP. You can specify up or down. If you specify the direction as up, CCMs are transmitted out of every logical interface that is part of the same bridging or VPLS instance except for the interface configured on the MEP. If you specify the direction as down, CCMs are transmitted only out of the interface configured on the MEP.

**NOTE:** Ports in the Spanning Tree Protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in an STP blocking state without the continuity check protocol configured do block CFM packets.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
user@host# set direction down
```

**NOTE:** Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 5G Universal Routing Platforms, you no longer need to configure the **no-control-word** statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. For all other interfaces on MX Series routers and on all other routers and switches, you must continue to configure the **no-control-word** statement at the **[edit routing-instances *routing-instance-name* protocols l2vpn]** or **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** hierarchy level when you configure CFM MEPs. Otherwise, the CFM packets are not transmitted, and the **show oam ethernet connectivity-fault-management mep-database** command does not display any remote MEPs.

15. Specify the interface to which the MEP is attached. It can be a physical interface, logical interface, or trunk interface. On MX Series routers, the MEP can be attached to a specific VLAN of a trunk interface. Also, specify the transport path as working.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name mep mep-id]
user@host# set interface interface-name protect
```

## SEE ALSO

[auto-discovery | 999](#)

[interval | 1038](#)

[name-format | 1058](#)

[protect-maintenance-association | 1070](#)

[short-name-format | 1076](#)

## Configuring Linktrace Protocol in CFM

The linktrace protocol is used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the **traceroute** command to verify the path between a pair of MEPs under the same maintenance association. Linktrace messages can also be used to verify the path between an MEP and an MIP under the same maintenance domain. The linktrace protocol enables you to configure the time to wait for a response. If no response is received for a linktrace request message, the request and response entries are deleted after the interval expires. You can also configure the number of linktrace reply entries to be stored for the corresponding linktrace request.

The operation of IEEE 802.1ag linktrace request and response messages is similar to the operation of Layer 3 **traceroute** commands. For more information about the **traceroute** command, see the *Junos OS Administration Library*.

To configure the linktrace protocol:

1. Configure the time to wait for a linktrace response at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level. You can specify the value in minutes or seconds. The default value is 10 minutes.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set linktrace age time
```

2. Configure the number of linktrace reply entries to be stored per linktrace request. You can specify a value from 1 through 500. The default value is 100.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set linktrace path-database-size path-database-size
```

### SEE ALSO

[age | 1102](#)

[path-database-size | 1060](#)

[connectivity-fault-management | 1132](#)

## Configuring Rate Limiting of Ethernet OAM Messages

The M320 with Enhanced III FPC, M120, M7i, M10 with CFEB, and MX Series routers support rate limiting of Ethernet OAM messages. Depending on the connectivity fault management (CFM) configuration, CFM

packets are discarded, sent to the CPU for processing, or flooded to other bridge interfaces. This feature allows the router to intercept incoming CFM packets for prevention of DoS attacks.

You can apply rate limiting of Ethernet OAM messages at either of two CFM policing levels, as follows:

- Global-level CFM policing—uses a policer at the global level to police the CFM traffic belonging to all the sessions.
- Session-level CFM policing—uses a policer created to police the CFM traffic belonging to one session.

To configure global-level CFM policing, include the **policer** statement and its options at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

To configure session-level CFM policing, include the **policer** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* level *number* maintenance-association *ma-name*]** hierarchy level.

The following example shows a CFM policer used for rate-limiting CFM:

```
[edit]
firewall {
  policer cfm-policer {
    if-exceeding {
      bandwidth-limit 8k;
      burst-size-limit 2k;
    }
    then discard;
  }
}
```

### Case 1: Global-Level CFM Policing

This example shows a global level policer, at the CFM level, for rate-limiting CFM. The **continuity-check *cfm-policer*** statement at the global **[edit protocols oam ethernet connectivity-fault-management policer]** hierarchy level specifies the policer to use for policing all continuity check packets of the CFM traffic belonging to all sessions. The **other *cfm-policer1*** statement at the **[edit protocols oam ethernet connectivity-fault-management policer]** hierarchy level specifies the policer to use for policing all non-continuity check packets of the CFM traffic belonging to all sessions. The **all *cfm-policer2*** statement specifies to police all CFM packets with the specified policer *cfm-policer2*. If the **all *policer-name*** option is used, then the user cannot specify the previous **continuity-check** and **other** options.

```
[edit protocols oam ethernet]
connectivity-fault-management {
```

```

    policer {
        continuity-check cfm-policer;
        other cfm-policer1 ;
        all cfm-policer2;
    }
}

```

## Case 2: Session-Level CFM Policing

This example shows a session-level CFM policer used for rate-limiting CFM. The **policer** statement at the session **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]** hierarchy level specifies the policer to use for policing only continuity check packets of the CFM traffic belonging to the specified session. The **other *cfm-policer1*** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name*]** hierarchy level specifies the policer to use for policing all non-continuity check packets of the CFM traffic belonging to this session only. The **all *cfm-policer2*** statement specifies to police all CFM packets with the specified policer *cfm-policer2*. If the **all *policer-name*** option is used, then the user cannot specify the previous **continuity-check** and **other** options.

```

[edit protocols oam ethernet]
connectivity-fault-management {
    maintenance-domain md {
        level number;
        maintenance-association ma {
            continuity-check {
                interval 1s;
            }
            policer {
                continuity-check cfm-policer;
                other cfm-policer1;
                all cfm-policer2;
            }
        }
        mep 1 {
            interface ge-3/3/0.0;
            direction up;
            auto-discovery;
        }
    }
}

```

```
}
```

In the case of global CFM policing, the same policer is shared across multiple CFM sessions. In per-session CFM policing, a separate policer must be created to rate-limit packets specific to that session.

**NOTE:**

Service-level policer configuration for any two CFM sessions on the same interface at different levels must satisfy the following constraints if the direction of the sessions is the same:

- If one session is configured with **policer all**, then the other session cannot have a **policer all** or **policer other** configuration.
- If one session is configured with **policer other**, then the other session cannot have a **policer all** or **policer other** configuration.

A commit error will occur if such a configuration is committed.

**NOTE:** Policers with PBB and MIPs are not supported.

SEE ALSO

[policer | 1064](#)

[policer | 1063](#)

[show oam ethernet connectivity-fault-management policer | 1934](#)

[clear oam ethernet connectivity-fault-management policer | 1442](#)

## Configuring Ethernet Local Management Interface

### IN THIS SECTION

- [Ethernet Local Management Interface Overview | 580](#)
- [Configuring the Ethernet Local Management Interface | 582](#)
- [Example E-LMI Configuration | 584](#)

### Ethernet Local Management Interface Overview


Gigabit Ethernet (**ge**), 10-Gigabit Ethernet (**xe**), and Aggregated Ethernet (**ae**) interfaces support the 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.

The E-LMI specification is available at the Metro Ethernet Forum. E-LMI procedures and protocols are used for enabling automatic configuration of the customer edge (CE) to support Metro Ethernet services. The E-LMI protocol also provides user-to-network interface (UNI) and Ethernet virtual connection (EVC) status information to the CE. The UNI and EVC information enables automatic configuration of CE operation based on the Metro Ethernet configuration.

The E-LMI protocol operates between the CE device and the provider edge (PE) device. It runs only on the PE-CE link and notifies the CE of connectivity status and configuration parameters of Ethernet services available on the CE port. The scope of the E-LMI protocol is shown in [Figure 10 on page 580](#).

**Figure 10: Scope of the E-LMI Protocol**



The E-LMI implementation on ACX and MX Series routers includes only the PE side of the E-LMI protocol.

E-LMI interoperates with an OAM protocol, such as Connectivity Fault Management (CFM), that runs within the provider network to collect OAM status. CFM runs at the provider maintenance level (UNI-N to UNI-N with up MEPs at the UNI). E-LMI relies on the CFM for end-to-end status of EVCs across CFM domains (SVLAN domain or VPLS).

The E-LMI protocol relays the following information:

- Notification to the CE of the addition/deletion of an EVC (active, not active, or partially active)
- Notification to the CE of the availability state of a configured EVC
- Communication of UNI and EVC attributes to the CE:
  - UNI attributes:
    - UNI identifier (a user-configured name for UNI)
    - CE-VLAN ID/EVC map type (all-to-one bundling, service multiplexing with bundling, or no bundling)



- Bandwidth profile is not supported (including the following features):
  - CM (coupling mode)
  - CF (color flag)
  - CIR (committed Information rate)
  - CBR (committed burst size)
  - EIR (excess information rate)
  - EBS (excess burst size)
- EVC attributes:
  - EVC reference ID
  - EVC status type (active, not active, or partially active)
  - EVC type (point-to-point or multipoint-to-multipoint)
  - EVC ID (a user-configured name for EVC)
  - Bandwidth profile (not supported)
- CE-VLAN ID/EVC map

E-LMI on MX Series routers supports the following EVC types:

- Q-in-Q SVLAN (point-to-point or multipoint-to-multipoint)—Requires an end-to-end CFM session between UNI-Ns to monitor the EVS status.
- VPLS (BGP or LDP) (point-to-point or multipoint-to-multipoint)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.
- L2 circuit/L2VPN (point-to-point)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.

**NOTE:** **l2-circuit** and **l2vpn** are not supported.

The E-LMI protocol on ACX Series routers supports Layer 2 circuit and Layer 2 VPN EVC types and enables link-loss forwarding for pseudowire (Layer 2 circuit and Layer 2 VPN) services as follows:

- Interworking between the connectivity fault management (CFM) protocol and the E-LMI protocol for Layer 2 circuit and Layer 2 VPN.
  - End-to-end CFM session between UNIs to monitor EVC status.
  - In the case of pseudowire redundancy, CFM can be used to monitor active and backup pseudowire sessions. The EVC status is declared as down to CE devices only when both the active and backup pseudowire sessions go down.

- Interworking between remote defect indication (RDI) and E-LMI for Layer 2 circuit and Layer 2 VPN.
  - If a maintenance association end point (MEP) receives an RDI bit set in a continuity check message (CCM) frame, and if RDI fault detection is enabled in the EVC configuration at **[edit protocols oam ethernet evcs evc-id evc-protocol cfm management-domain name management-association name faults rdi]**, then the pseudowire is declared as down to CE routers through E-LMI.
- If an end-to-end CFM session does not exist between UNIs, the pseudowire (Layer 2 circuit or Layer 2 VPN) up and down state triggers an asynchronous EVC state change message to CE routers through E-LMI.

**NOTE:** ACX Series routers do not support E-LMI for Layer 2 services (bridging).

### Configuring the Ethernet Local Management Interface

#### IN THIS SECTION

- [Configuring an OAM Protocol \(CFM\) | 582](#)
- [Assigning the OAM Protocol to an EVC | 582](#)
- [Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC | 583](#)

To configure E-LMI, perform the following steps:

#### Configuring an OAM Protocol (CFM)

For information on configuring the OAM protocol (CFM), see [“IEEE 802.1ag OAM Connectivity Fault Management Overview” on page 548](#).

#### Assigning the OAM Protocol to an EVC

To configure an EVC, you must specify a name for the EVC using the **evcsevc-id** statement at the **[edit protocols oam ethernet]** hierarchy level. You can set the EVC protocol for monitoring EVC statistics to **cfm** or **vpls** using the **evc-protocol** statement and its options at the **[edit protocols oam ethernet evcs]** hierarchy level.

You can set the number of remote UNIs in the EVC using the **remote-uni-count number** statement at the **[edit protocols oam ethernet evcs evcs-protocol]** hierarchy level. The **remote-uni-count** defaults to 1. Configuring a value greater than 1 makes the EVC multipoint-to-multipoint. If you enter a value greater than the actual number of endpoints, the EVC status will display as partially active even if all endpoints are up. If you enter a **remote-uni-count** less than the actual number of endpoints, the status will display as active, even if all endpoints are not up.

You can configure an EVC by including the **evcs** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
evcs evc-id {
  evc-protocol (cfm (management-domain name management-association name ) | vpls (routing-instance name))
  {
    remote-uni-count <number>; # Optional, defaults to 1
    multipoint-to-multipoint;
    # Optional, defaults to point-to-point if remote-uni-count is 1
  }
}
```

### **Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC**

To configure E-LMI, include the **lmi** statement at the **[edit protocols oam ethernet]** hierarchy level:

```
[edit protocols oam ethernet]
lmi {
  polling-verification-timer value;
  # Polling verification timer (T392), defaults to 15 seconds
  status-counter count; # Status counter (N393), defaults to 4
  interface name {
    evc evc-id {
      default-evc;
      vlan-list [ vlan-ids ];
    }
    evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
    polling-verification-time value; # Optional, defaults to global value
    status-counter count; # Optional, defaults to global value
    uni-id value; # Optional, defaults to interface-name
  }
}
```

You can set the status counter to count consecutive errors using the **status-counter count** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The status counter is used to determine if E-LMI is operational or not. The default value is 4.

You can set the **polling-verification-timer value** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. The default value is 15 seconds.

You can enable an interface and set its options for use with E-LMI using the **interface name** statement at the **[edit protocols oam ethernet lmi]** hierarchy level. Only **ge**, **xe**, and **ae** interfaces are supported. You can use the interface **uni-id** option to specify a name for the UNI. If **uni-id** is not configured, it defaults to the name variable of **interface name**.

You can specify the CE-VLAN ID/EVC map type using the **evc-map-type type** interface option. The options are **all-to-one-bundling**, **bundling**, or **service-multiplexing**. Service multiplexing is with no bundling. The default type is **all-to-one-bundling**.

To specify the EVC that an interface uses, use the **evc evc-id** statement at the **[edit protocols oam ethernet lmi interface name]** hierarchy level. You can specify an interface as the default EVC interface using the **default-evc** statement at the **[edit protocols oam ethernet lmi interface name evc evc-id]** hierarchy level. All VLANs that are not mapped to any other EVCs are mapped to this EVC. Only one EVC can be configured as the default.

You can map a list of VLANs to an EVC using the **vlan-list vlan-id-list** statement at the **[edit protocols oam ethernet lmi interface name evc evc-id]** hierarchy level.

### Example E-LMI Configuration

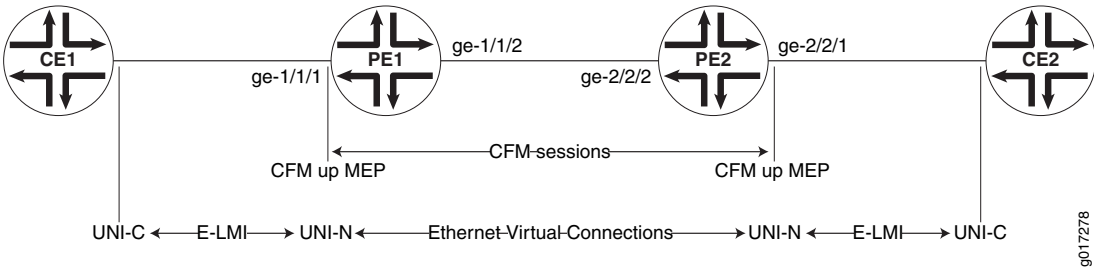
#### IN THIS SECTION

- [Example Topology | 584](#)
- [Configuring PE1 | 584](#)
- [Configuring PE2 | 586](#)
- [Configuring Two UNIs Sharing the Same EVC | 588](#)

### Example Topology

Figure 11 on page 584 illustrates the E-LMI configuration for a point-to-point EVC (SVLAN) monitored by CFM. In this example, VLANs 1 through 2048 are mapped to **evc1** (SVLAN 100) and 2049 through 4096 are mapped to **evc2** (SVLAN 200). Two CFM sessions are created to monitor these EVCs.

Figure 11: E-LMI Configuration for a Point-to-Point EVC (SVLAN) Monitored by CFM



### Configuring PE1

```
[edit]
interfaces {
```

```

ge-1/1/1 {
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-2048;
    }
  }
  unit 1 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 2049-4096;
    }
  }
}
ge-1/1/2 {
  unit 0 {
    vlan-id 100;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 1-2048;
    }
  }
  unit 1 {
    vlan-id 200;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 2049-4096;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain md {
          level 0;
          maintenance-association 1 {
            name-format vlan;
            mep 1 {
              direction up;
              interface ge-1/1/1.0 vlan 1;
            }
          }
        }
      }
    }
  }
}

```

```

        maintenance-association 2049 {
            name-format vlan;
            mep 1 {
                direction up;
                interface ge-1/1/1.1 vlan 2049;
            }
        }
    }
}
evcs {
    evc1 {
        evc-protocol cfm management-domain md management-association 1;
        remote-uni-count 1;
    }
    evc2 {
        evc-protocol cfm management-domain md management-association 2049;
        remote-uni-count 1;
    }
}
lmi {
    interface ge-1/1/1 {
        evc evc1 {
            vlan-list 1-2048;
        }
        evc evc2 {
            vlan-list 2049-4096;
        }
        evc-map-type bundling;
        uni-id uni-ce1;
    }
}
}
}
}

```

### Configuring PE2

```

[edit]
interfaces {
    ge-2/2/1 {
        unit 0 {
            family bridge {
                interface-mode trunk;
                vlan-id-list 1-2048;
            }
        }
    }
}

```

```

    }
  }
  unit 1 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 2049-4096;
    }
  }
}
ge-2/2/2 {
  unit 0 {
    vlan-id 100;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 1-2048;
    }
  }
  unit 1 {
    vlan-id 200;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 2049-4095;
    }
  }
}
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain md {
          level 0;
          maintenance-association 1 {
            name-format vlan;
            mep 1 {
              direction up;
              interface ge-2/2/1.0 vlan 1;
            }
          }
          maintenance-association 2049 {
            name-format vlan;
            mep 1 {
              direction up;
              interface ge-2/2/1.1 vlan 2049;
            }
          }
        }
      }
    }
  }
}

```

```

    }
  }
}
}
evcs {
  evc1 {
    evc-protocol cfm management-domain md management-association 1;
    remote-uni-count 1;
  }
  evc2 {
    evc-protocol cfm management-domain md management-association 2049;
    uni-count 2;
  }
}
lmi {
  interface ge-2/2/1 {
    evc evc1 {
      vlan-list 1-2048;
    }
    evc evc2 {
      vlan-list 2049-4095;
    }
    evc-map-type bundling;
    uni-id uni-ce2;
  }
}
}
}
}

```

### ***Configuring Two UNIs Sharing the Same EVC***

```

[edit protocols]
oam {
  ethernet {
    connectivity-fault-management { ...}
    evcs {
      evc1 {
        evc-protocol cfm management-domain md management-association 1;
        remote-uni-count 1;
      }
    }
  }
  lmi {
    interface ge-2/2/1 {

```



```

    evc evc1 {
        vlan-list 0-4095;
    }
    evc-map-type all-to-one-bundling;
    uni-id uni-ce1;
}
interface ge-2/3/1 {
    evc evc1 {
        vlan-list 0-4095;
    }
    evc-map-type all-to-one-bundling;
    uni-id uni-ce2;
}
}
}
}

```

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | 1132

### Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

You can create a connectivity fault management (CFM) action profile to define event flags and thresholds to be monitored. You can also specify the action to be taken when any of the configured events occur. When the CFM events occur, the router performs the corresponding action based on your specification. You can configure one or more events in the action profile. Alternatively, you can configure an action profile and specify default actions when connectivity to a remote maintenance association endpoint (MEP) fails.

**NOTE:** You cannot configure multiple actions at this time. Only one action can be configured. This limitation affects both the **action** and **clear-action** statements.

To configure the CFM action profile:

1. In configuration mode, at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level, specify the name of the action profile and the CFM event(s). You can configure more than one event in the action profile. Possible events include: interface-status-tlv, port-status-tlv, adjacency-loss, RDI.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set action-profile profile-name event [event1, event2, event3..]
```

2. Specify the action to be taken by the router when the event occurs. The action is triggered when the event occurs. If you have configured more than one event in the action profile, it is not necessary for all events to occur to trigger the action.

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

3. Specify the default action to be taken by the router when connectivity to a remote MEP fails. If no action is configured, no action is taken.

**NOTE:** Associating an action profile with the **interface-down** action on an up MEP CFM session running over a circuit cross-connect (CCC) interface (l2circuit/l2vpn) is not advisable and can result in a deadlock situation.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set action-profile profile-name default-actions action
```

## SEE ALSO

---

[event \(CFM\) | 1025](#)

---

[default-actions | 1016](#)

---

[connectivity-fault-management | 1132](#)


## CFM Action Profile to Bring Down a Group of Logical Interfaces Overview

With growing networks, there is a requirement of monitoring a large number of services using CFM. To monitor each service, one session per service logical interface is required. If the services are large in number, this method does not scale as the number of sessions are limited. Instead of one CFM session per service, a single CFM session can monitor multiple services.

Also, there are scenarios where the user-to-network interface (UNI) device needs to be brought down based on sessions on network-to-network Interface (NNI) logical interface. Here, the NNI logical interface refers to core interface and UNI physical interface refers to access interface hosting multiple service logical interfaces. Based on core interface monitoring, you can bring down service logical interfaces associated with access interface.

Figure 12 on page 591 illustrates a topology where a number of services destined to customer-edge (CE) routers share a single port on a provider-edge (PE) router. Each service uses one logical interface. A set of services or logical interfaces (colored in yellow) are destined to one CE router and a set of services or logical interfaces colored in red are destined to another CE router. To monitor each service, you need dedicated down maintenance association end point (MEP) sessions for each service. You can bring down the service by bringing down the service logical interface whenever the session goes down. However, this approach is not scalable if we have large number of services. Monitoring the CFM session on the physical interface is also not feasible because multiple CE routers might be connected and the services to other CE router could be disrupted. To address this issue of monitoring multiple services with a single session, you can create a CCM action profile to bring down a group of logical interfaces by using a CFM session that is configured on a single logical interface.

**Figure 12: Topology of Multiple VLAN Services Sharing a Single Port on PE Router Destined to Multiple CE Routers**



You can configure CCM action profiles for the following scenarios:

- To bring down a group of logical interfaces all having the same parent port when CCM monitoring session is running on one of the logical interface but on a different parent port.
- To bring down a group of logical interfaces when CCM monitoring session is running on one of the logical interfaces, all belonging to the same parent port.
- To bring down the port, when the CCM monitoring session is running on one of the logical interfaces of a different parent port.

### **Benefits of Creating CFM Action Profile to Bring Down a Group of Logical Interfaces**

- Reduces resource requirement in scaled networks where multiple services need to be monitored.
- Avoids the need to create individual MEP sessions for each service in a topology that includes multiple services to be monitored, thereby enhancing the performance and scalability of the network.

SEE ALSO

[action-profile](#) | 996

### **Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces**

To monitor multiple services or IFLs using CFM session configured on a single logical interface, you can create a CCM action profile to bring down a group of logical interfaces. You need to define an action to bring down the interface group in the action profile. You will then define the interface device name and the number of logical interfaces that have to be brought down. A logical interface is represented by a combination of the **interface-device-name** and **unit-list**. The following steps explain the procedure to bring down a group of logical interfaces when the **interface-device-name** and/or **unit-list** are specified.

1. In configuration mode, at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level, specify the name of the action profile and the CFM event(s). You can configure more than one event in the action profile.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set action-profile profile-name event [event1, event2, event3..]
```

For example,

```
user@host# set action-profile AP_test event adjacency-loss rdi
```

**NOTE:** The action **interface-group-down** will not be supported with events other than adjacency-loss and RDI. Any other events configured results in a commit error.

2. In configuration mode, at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name*]** hierarchy level, define the action to bring down the interface group.

```
[edit protocols oam ethernet connectivity-fault-management action-profile AP-test ]
user@host# set action interface-group-down
```

**NOTE:** The action **interface-group-down** will not be supported with other interface related actions. Any other actions configured results in a commit error.

3. At the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level, define the maintenance domain. Specify the maintenance-association parameters.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# set maintenance-domain domain-name level number maintenance-association ma-name
continuity-check interval 1s
```

For example,

```
user@host# set maintenance-domain md6 level 6 maintenance-association ma6 continuity-check interval 1s
```

4. At the **edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name***, define the maintenance association endpoint and the associated parameters.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name]
user@host# set mep mep-id interface interface-name direction down remote -mep mep-id
```

For example,

```
user@host# set mep 101 interface ge-0/0/0.0 direction down remote -mep 102
```

5. If the action-profile has **interface-group-down** action configured, it is mandatory to configure the **interface-group** at the RMEP level. In the configuration mode at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id* action-profile *profile-name*]** include the **interface-group** statement to bring down the interface group marked with the action profile as **interface-group-down**.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep mep-id action-profile profile-name]
user@host# set interface-group
```

For example,

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md6
  maintenance-association ma6 mep 101 remote-mep 102 action-profile AP_test]
user@host# set interface-group
```

**NOTE:** If the **interface-group** configuration is not included in the RMEP configuration. The configuration results in commit error.

6. A logical interface is represented by a combination of the **interface-device-name** and **unit-list**. Configure the device interface name and the number of logical interfaces at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id* action-profile *profile-name* interface-group**.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id remote-mep mep-id action-profile profile-name
  interface-group]
user@host# set interface interface-name
user@host# set unit-list logical-interface-unit-number
```

For example,

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md6
  maintenance-association ma6 mep 101 remote-mep 102 action-profile AP_test interface-group]
user@host# set interface ge-0/0/0.0
user@host# set unit-list 1223-3344
```

In this configuration example, the interface ge-0/0/0.0 is brought down.

**NOTE:**

- At least one of the **interface-group** parameters, **interface-device-name** or **unit-list** must be configured. If the interface device name is not configured, the MEP interface is considered as the device name and the logical interface on that device is brought down.
- If the **unit-list** parameter exceeds the recommended limit, a commit error occurs.
- If the *interface-device-name* is not specified in the **interface-group**, the logical interface numbers mentioned in **unit-list** for the physical interface is brought down.
- If the *unit-list* is not specified in the **interface-group**, IFLs are brought down for the configured interface.

7. Verify the configuration using **show protocols oam** command.

```
[edit]
user@host# show protocols oam
ethernet {
  connectivity-fault-management {
    action-profile AP_TEST {
      event {
        adjacency-loss;
        rdi;
      }
      action {
        interface-group-down;
      }
    }
  }
  maintenance-domain md6 {
    level 6;
    maintenance-association ma6 {
      continuity-check {
        interval 1s;
      }
      mep 102 {
        interface ge-0/0/0.0;
        direction down;
        remote-mep 103 {
          action-profile AP_TEST;
          interface-group {
            ge-0/0/1;
            unit-list [12 23-33 44];
          }
        }
      }
    }
  }
}
}
```

#### SEE ALSO

[interface-group | 1214](#)

[interface-group-down | 1215](#)

## Enabling Enhanced Connectivity Fault Management Mode

You can enable enhanced connectivity fault management (CFM) mode to enable effective Ethernet OAM deployment in scaling networks. On enabling enhanced CFM mode, Junos OS supports 32,000 maintenance association end points (MEPs) and maintenance intermediate points (MIPs) each per chassis for bridge, VPLS, L2VPN, and CCC domains. In previous releases, Junos OS supports 8,000 MEPs and 8000 MIPS per chassis. If you do not enable enhanced CFM, Junos OS continues to support existing number of MIPs and MEPs per chassis.

**NOTE:** To support enhanced CFM mode, configure the network services mode on the router as **enhanced-ip**. If the network services mode is not **enhanced-ip**, and you have enabled enhanced CFM, the following warning message is displayed:

```
[edit protocols oam ethernet]
'connectivity-fault-management'
enhanced ip is not effective please configure enhanced ip and give router reboot
```

To enable enhanced CFM mode, perform the following steps:

1. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management
```

2. Enable effective Ethernet OAM deployment by enabling enhanced CFM mode.

```
[edit protocols oam ethernet connectivity-fault-management ]
user@host# set enhanced-cfm-mode
```

3. Commit the mode change. A warning message is displayed asking you to restart CFM. If you do not restart CFM, CFM is automatically restarted by Junos OS.

```
[edit protocols oam ethernet connectivity-fault-management ]
user@host # commit
[edit protocols oam ethernet]
'connectivity fault management'
CFM mode change is catastrophic. cfmd will be restarted
commit complete
```



4. To verify if the enhanced CFM mode has been configured, use the **show oam ethernet connectivity-fault-management state** command.

```
[edit protocols oam ethernet connectivity-fault-management ]
user@host# run show oam ethernet connectivity-fault-management
enhanced-cfm-mode;
traceoptions {
    file cfmd.log size 1g;
}
maintenance-domain md6 {
    level 6;
    maintenance-association ma6 {
        continuity-check {
            interval 1s;
        }
        mep 102 {
            interface ge-0/0/0.0;
            direction up;
        }
    }
}
```

SEE ALSO

| [enhanced-cfm-mode](#) | [1023](#)

## Configuring M120 and MX Series Routers for CCC Encapsulated Packets

### IN THIS SECTION

- [IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets Overview](#) | [597](#)
- [CFM Features Supported on Layer 2 VPN Circuits](#) | [598](#)
- [Configuring CFM for CCC Encapsulated Packets](#) | [598](#)

### *IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets Overview*

Layer 2 virtual private network (L2VPN) is a type of virtual private network service used to transport customer's private Layer 2 traffic (for example, Ethernet, ATM or Frame Relay) over the service provider's

shared IP/MPLS infrastructure. The service provider edge (PE) router must have an interface with circuit cross-connect (CCC) encapsulation to switch the customer edge (CE) traffic to the public network.


The IEEE 802.1ag Ethernet Connectivity Fault Management (CFM) is an OAM standard used to perform fault detection, isolation, and verification on virtual bridge LANs. M120 and MX Series routers provide CFM support for bridge/VPLS/routed interfaces and support 802.1ag Ethernet OAM for CCC encapsulated packets.

### CFM Features Supported on Layer 2 VPN Circuits

CFM features supported on L2VPN circuits are as follows:

- Creation of up/down MEPs at any level on the CE-facing logical interfaces.
- Creation of MIPs at any level on the CE-facing logical interfaces.
- Support for continuity check, loopback, and linktrace protocol.
- Support for the Y1731 Ethernet Delay measurement protocol.
- Support for action profiles to bring the CE-facing logical interfaces down when loss of connectivity is detected.

Figure 13: Layer 2 VPN Topology



To monitor the L2VPN circuit, a CFM up MEP (Level 6 in [Figure 13 on page 598](#)) can be configured on the CE-facing logical interfaces of provider edge routers PE1 and PE2. To monitor the CE-PE attachment circuit, a CFM down MEP can be configured on the customer logical interfaces of CE1-PE1 and CE2-PE2 (Level 0 in [Figure 13 on page 598](#)).

### Configuring CFM for CCC Encapsulated Packets

The only change from the existing CLI configuration is the introduction of a new command to create a MIP on the CE-facing interface of the PE router.

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        # Define a maintenance domains for each default level.
        #; These names are specified as DEFAULT_level_number
        maintenance-domain DEFAULT_x {
```

```

# L2VPN CE interface
interface (ge | xe)-fpc/pic/port.domain;
}
{
level number;
maintenance-association identifier {
    mep mep-id {
        direction (up | down);
        # L2 VPN CE interface on which encapsulation family CCC is configured.
        interface (ge | xe)-fpc/pic/port.domain;
        auto-discovery;
        priority number;
    }
}
}
}
}
}
}
}
}

```

SEE ALSO

| [connectivity-fault-management](#) | 1132

## Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades

Starting in Release 17.1, Junos OS connectivity fault management (CFM), during a unified in-service software upgrade (ISSU), works when the peer device is not a Juniper Networks router. Interoperating with the router of another vendor, the Juniper Networks router retains session information and continues to transmit continuity check message (CCM) PDUs during the unified ISSU. Connectivity fault management continues to operate.

This feature requires the following conditions be met:

- Packet Forwarding Engine keepalives must be enabled to provide inline transmission of CCMs. The feature does not work when the CCMs are transmitted by the CPU of a line card, which is the default transmission method.
- The interval between CCMs must be 1 second.

CFM interoperability during a unified ISSU is supported on the following MPCs: MPC1, MPC2, MPC2-NG, MPC3-NG, MPC5, and MPC6.

To enable CFM interoperability with third-party devices across a unified ISSU:

1. Enable inline keepalives.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set hardware-assisted-keepalives enable
```

2. Set the CCM interval to 1 second.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name continuity-check]
user@host# set interval 1s
```

SEE ALSO

[Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling](#) | 825

## Configuring Unified ISSU for 802.1ag CFM

A unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is automatically enabled for the Connectivity Fault Management (CFM) protocols and interoperates between local and remote maintenance endpoints (MEPs).

The Junos OS provides support for unified ISSU using the loss threshold type length value (TLV), which is automatically enabled for CFM. TLVs are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and optional information in a protocol data unit (PDU). The loss threshold TLV indicates the loss threshold value of a remote MEP. The loss threshold TLV is transmitted as part of the CFM continuity check messages.

**NOTE:** Starting in Junos OS Release 15.1, configuring ISSU with CFM (802.1ag) is supported only on MX and PTX routers that support TLV. Interoperation with other vendors is not supported.

During a unified ISSU, the control plane may go down for several seconds and cause CFM continuity check packets to get dropped. This may cause the remote MEP to detect a connectivity loss and mark the MEP as down. To keep the MEP active during a unified ISSU, the loss threshold TLV communicates the minimum threshold value the receiving MEP requires to keep the MEP active. The receiving MEP parses the TLV and updates the loss threshold value, but only if the new threshold value is greater than the locally configured threshold value.

An overview of CFM is described starting in [“IEEE 802.1ag OAM Connectivity Fault Management Overview” on page 548](#), and you should further observe the additional requirements described in this topic.

[Table 84 on page 601](#) shows the Loss Threshold TLV format.

**Table 84: Loss Threshold TLV Format**

| Parameter | Octet (sequence) | Description                                                                                                                                                                                                              |
|-----------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type=31   | 1                | Required. Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.                                                                                                  |
| Length=12 | 2                | Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field. |
| OUI       | 3                | Optional. Organization unique identifier (OUI), which is controlled by the IEEE and is typically the first three bytes of a MAC address (Juniper OUI 0x009069).                                                          |
| Subtype   | 1                | Optional. Organizationally defined subtype.                                                                                                                                                                              |
| Value     | 4                | Optional. Loss threshold value.                                                                                                                                                                                          |
| Flag      | 4                | Optional. Bit0 (identifies an ISSU is in progress)<br><br>Bit1-31 (reserved)                                                                                                                                             |

Junos OS provides configuration support for the **convey-loss-threshold** statement, allowing you to control the transmission of the loss threshold TLV in continuity check messages PDUs. The **convey-loss-threshold** statement specifies that the loss threshold TLV must be transmitted as part of the continuity check messages. If the **convey-loss-threshold** statement is not specified, continuity check messages transmit this TLV only when a unified ISSU is in progress. The Junos OS provides this configuration at the continuity-check level. By default, continuity check messages do not include the loss threshold TLV.

To configure the convey loss threshold, use the **convey-loss-threshold** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain identifier maintenance-association identifier continuity-check]** hierarchy level.

For the remote MEP, the loss threshold TLV is transmitted only during the unified ISSU if the **convey-loss-threshold** statement is not configured. The remote MEP switches back to the default loss threshold if no loss threshold TLV is received or the TLV has a default threshold value of 3.

An example of the ISSU configuration statements follows:

```
protocols {
```

```

oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain identifier {
        level number;
        maintenance-association identifier {
          continuity-check {
            convey-loss-threshold;
            interval number;
            loss-threshold number;
            hold-interval number;
          }
        }
      }
    }
  }
}

```

The Junos OS saves the last received loss threshold TLV from the remote MEP. You can display the last saved loss threshold TLV that is received by the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```

user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md3 maintenance-association ma5 local-mep 2 remote-mep 1

```

```

Maintenance domain name: md3, Format: string, Level: 3
Maintenance association name: ma3, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: yes
  Prefer me: no, Protection in use: no, FRR Flag: no
Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
Loss Threshold TLV:
  Loss Threshold: 3 , Flag: 0x0

Remote MEP identifier: 1, State: ok
MAC address: 00:1f:12:b7:ce:79, Type: Learned
Interface: xe-4/1/1.0
Last flapped: Never

```

```

Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
Effective loss threshold: 3 frames
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
Connection Protection TLV:
  Prefer me: no, Protection in use: no, FRR Flag: no
Loss Threshold TLV:   #Displays last received value
  Loss Threshold: 3   , Flag:      0x0

```

The Junos OS saves the last transmitted loss threshold TLV from a local MEP. You can display the last transmitted loss threshold TLV and the effective loss (operational) threshold for the remote MEP, using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```

user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md3 maintenance-association ma5 local-mep 2 remote-mep 1

```

```

Maintenance domain name: md3, Format: string, Level: 3
Maintenance association name: ma3, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 2, Direction: up, MAC address: 00:19:e2:b0:76:be
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: yes
  Prefer me: no, Protection in use: no, FRR Flag: no
Interface name: xe-4/1/1.0, Interface status: Active, Link status: Up
Loss Threshold TLV:   #Displays last transmitted value
  Loss Threshold: 3   , Flag:      0x0

Remote MEP identifier: 1, State: ok
MAC address: 00:1f:12:b7:ce:79, Type: Learned
Interface: xe-4/1/1.0
Last flapped: Never
Continuity: 100%, Admin-enable duration: 45sec, Oper-down duration: 0sec
Effective loss threshold: 3 frames   #Displays operational threshold
Remote defect indication: falsePort status TLV: none
Interface status TLV: none
Connection Protection TLV:
  Prefer me: no, Protection in use: no, FRR Flag: no

```

```

Loss Threshold TLV:
  Loss Threshold: 3 , Flag: 0x0

```

## SEE ALSO

*Before You Begin a Unified ISSU*

*Unified ISSU System Requirements*

## Junos OS Support for Performance Monitoring Compliant with Technical Specification MEF 36

Junos OS release 16.1R1 and later supports performance monitoring that is compliant with Technical Specification MEF 36. Technical Specification MEF 36 specifies the performance monitoring MIB. The performance monitoring MIB is required to manage service operations, administration, and maintenance (OAM) implementations that satisfy the Service OAM requirements and framework specified in MEF 17 and MEF 35, the management objects specified in MEF 7.1, and the performance monitoring functions defined in ITU-T Y.1731 and IEEE 802.1ag.

You can enable MEF-36-compliant performance monitoring by configuring the [measurement-interval](#) statement at the `[edit protocols oam ethernet cfm performance-monitoring]` hierarchy level.

When MEF-36-compliant performance monitoring is enabled:

- An SNMP get next request for a variable might not fetch the current value unless an SNMP walk is performed before performing the get next request. This limitation applies only to the current statistics for delay measurement, loss measurement, and synthetic loss measurement.
- The output for the field **Current delay measurement statistics** might display a measurement interval of 0 (zero) and an incorrect timestamp until the first cycle time has expired.
- Supported data TLV size for performance monitoring protocol data units (PDUs) is 1386 bytes when MEF-36-compliant performance monitoring is enabled. The TLV size is 1400 bytes in legacy mode.
- The maximum configurable value for the lower threshold bin is 4,294,967,294.
- Frame loss ratio (FLR) is excluded in loss measurements during period of unavailability for synthetic loss measurement only. In case of loss measurement, FLR is included even during period of unavailability.
- During a period of loss of continuity (adjacency down), although SOAM PDUs are not sent, FLR and availability calculations are not stopped. These calculations are performed with the assumption of 100% loss.



- The number of SOAM PDUs that are sent during the first measurement interval might be less than expected. This is because of a delay in detecting the adjacency state at the performance monitoring session level.
- The number of SOAM PDUs transmitted during a measurement interval for a cycle time of 100 ms might not be accurate. For example, in a measurement interval of two minutes with a cycle time 100 ms, the SOAM PDUs transmitted might be in the range of 1198–2000.

SEE ALSO

[measurement-interval](#) | 1053

## Damping CFM performance Monitoring Traps and Notifications to Prevent Congestion of The NMS

You can dampen the performance monitoring threshold-crossing traps and notifications that are generated every time a threshold-crossing event occurs to prevent congestion of the network management system (NMS).

Damping limits the number of `jnxSoamPmThresholdCrossingAlarm` traps sent to the NMS by summarizing the flap occurrences over a period of time, known as the flap trap timer, and sends a single `jnxSoamPmThresholdFlapAlarm` notification to the NMS. You can configure the duration of the flap trap timer to any value from 1 through 360 seconds.

The `jnxSoamPmThresholdFlapAlarm` notification is generated and sent when the following conditions are met:

- At least one flap has occurred when the flap timer has expired.
- You changed the value of the flap trap timer, which caused the timer to stop.

You can enable damping at the global level for the iterator or you can enable damping at the individual threshold type of the iterator. For instance, to enable damping at the global level, for the iterator, use the following command: **`set protocols oam ethernet cfm performance-monitoring sla-iterator-profiles profile-name flap-trap-monitor`**. To enable damping at a specific threshold type, for the `avg-fd-twoway-threshold`, use the following command: **`set protocols oam ethernet cfm performance-monitoring sla-iterator-profiles profile-name avg-fdv-twoway-threshold flap-trap-monitor`**.

You can also disable damping.

SEE ALSO

[flap-trap-monitor](#) | 1026

## Example: Configuring Ethernet CFM on Physical Interfaces

### IN THIS SECTION

- [Requirements | 606](#)
- [Overview | 606](#)
- [Configuration | 606](#)

This example shows the configuration of Ethernet connectivity fault management (CFM) on physical interfaces.

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.3 or later.

### Overview

CFM can be used to monitor the physical link between two routers. This functionality is similar to that supported by the IEEE 802.3ah LFM protocol.


In Junos OS Release 9.3 and later, CFM also supports aggregated Ethernet interfaces. On interfaces configured on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers, CFM is not supported on untagged aggregated Ethernet member links. MPCs and MICs do support CFM on untagged and tagged aggregated Ethernet logical interfaces.

**NOTE:** The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.

### Configuration

In the following example, two routers (Router 1 and Router 2) are connected by a point-to-point Gigabit Ethernet link. The link between these two routers is monitored using CFM. This is shown in [Figure 14 on page 607](#). The single boundary is a “down mep” in CFM terminology.

Figure 14: Ethernet CFM on Physical Interfaces



To configure Ethernet CFM on physical interfaces, perform these tasks:

### CLI Quick Configuration

#### Router 1

Configure the interface and CFM:

```
[edit]
interfaces ge-1/0/1 {
  unit 0 {
    family inet;
  }
}

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain private {
          level 0;
          maintenance-association private-ma {
            continuity-check {
              interval 1s;
            }
            mep 100 {
              interface ge-1/0/1;
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}
```

The configuration on Router 2 mirrors that on Router 1, with the exception of the *mep-id*.

## Router 2

Configure the interface and CFM:

```
[edit]
interfaces ge-0/2/5 {
  unit 0 {
    family inet;
  }
}

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain private {
          level 0;
          maintenance-association private-ma {
            continuity-check {
              interval 1s;
            }
            mep 200 {
              interface ge-0/2/5;
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}
```

To verify that the physical interface is configured correctly for CFM, use the **show interface** command. To verify the CFM configuration, use one or more of the **show oam ethernet connectivity-fault-management** commands listed in the [CLI Explorer](#).

## SEE ALSO

| [show oam ethernet connectivity-fault-management interfaces](#) | 1879

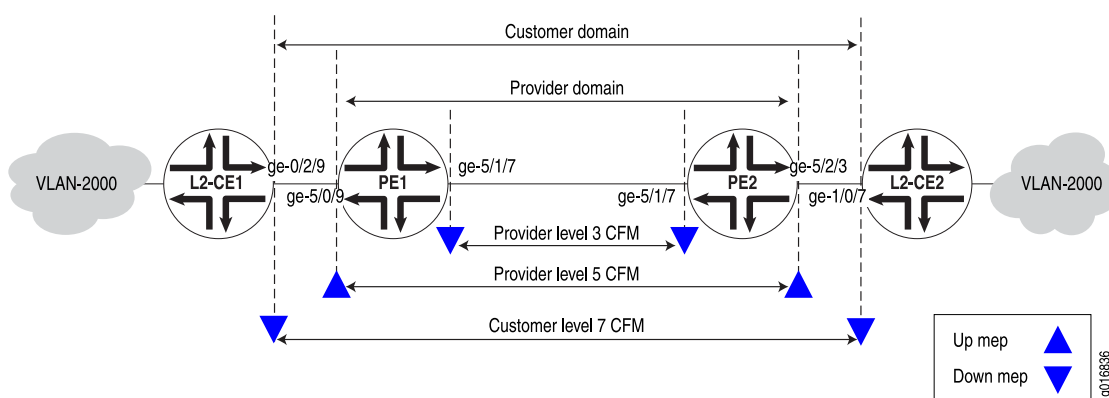
## Example: Configuring Ethernet CFM on Bridge Connections

In this example, both the customer and service provider are running Ethernet CFM over a simple bridge network. The network is shown in [Figure 15 on page 609](#). The customer has configured Ethernet CFM on MX Series routers L2-CE1 and L2-CE2. The service provider has configured Ethernet CFM on MX Series routers PE1 and PE2.

**NOTE:** The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.

The service provider is using CFM level 3 for the link between PE1 and PE2 and level 5 from one CE facing port to the other. The customer is using CFM level 7. The boundaries are marked with “up mep” and “down mep” CFM terminology in the figure.

Figure 15: Ethernet CFM over a Bridge Network



Here are the configurations of CFM on the customer routers.

### CFM on L2-CE1

```
[edit interfaces]
ge-0/2/9 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}
```

```
[edit protocols oam ethernet]
connectivity-fault-management {
  maintenance-domain customer {
    level 7;
    maintenance-association customer-site1 {
      continuity-check {
        interval 1s;
      }
      mep 700 {
        interface ge-0/2/9.0;
        direction down;
        auto-discovery;
      }
    }
  }
}
```

### CFM on L2-CE2

```
[edit interfaces]
ge-1/0/7 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam ethernet]
connectivity-fault-management {
  maintenance-domain customer {
    level 7;
    maintenance-association customer-site2 {
      continuity-check {
        interval 1s;
      }
      mep 800 {
        interface ge-1/0/7.0;
        direction down;
        auto-discovery;
      }
    }
  }
}
```

```

    }
  }
}

```

Here are the configurations of CFM on the provider routers.

### CFM on PE1

```

[edit interfaces]
ge-5/0/9 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 2000;
  }
}
ge-5/1/7 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 2000;
  }
}

[edit bridge-domains]
bridge-vlan2000 {
  domain-type bridge;
  vlan-id 2000;
  interface ge-5/0/9.0;
  interface ge-5/1/7.0;
}

[edit protocols oam ethernet connectivity-fault-management]
maintenance-domain provider-outer {
  level 5;
  maintenance-association provider-outer-site1 {
    continuity-check {
      interval 1s;

```

```

    }
    mep 200 {
        interface ge-5/0/9.0;
        direction up;
        auto-discovery;
    }
}
maintenance-domain provider-inner {
    level 3;
    maintenance-association provider-inner-site1 {
        continuity-check {
            interval 1s;
        }
        mep 200 {
            interface ge-5/1/7.0;
            direction down;
            auto-discovery;
        }
    }
}

```

## CFM on PE2

```

[edit interfaces]
ge-5/1/7 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}
ge-5/2/3 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 2000;
    }
}

```



```

}

[edit bridge-domains]
bridge-vlan2000 {
    domain-type bridge;
    interface ge-5/2/3.0;
    interface ge-5/1/7.0;
}

[edit protocols oam ethernet connectivity-fault-management]
maintenance-domain provider-outer {
    level 5;
    maintenance-association provider-outer-site1 {
        continuity-check {
            interval 1s;
        }
        mep 100 {
            interface ge-5/2/3.0;
            direction up;
            auto-discovery;
        }
    }
}
maintenance-domain provider-inner {
    level 3;
    maintenance-association provider-inner-site1 {
        continuity-check {
            interval 1s;
        }
        mep 100 {
            interface ge-5/1/7.0;
            direction down;
            auto-discovery;
        }
    }
}

```

SEE ALSO

| [Configuring Continuity Check Messages](#) | 648

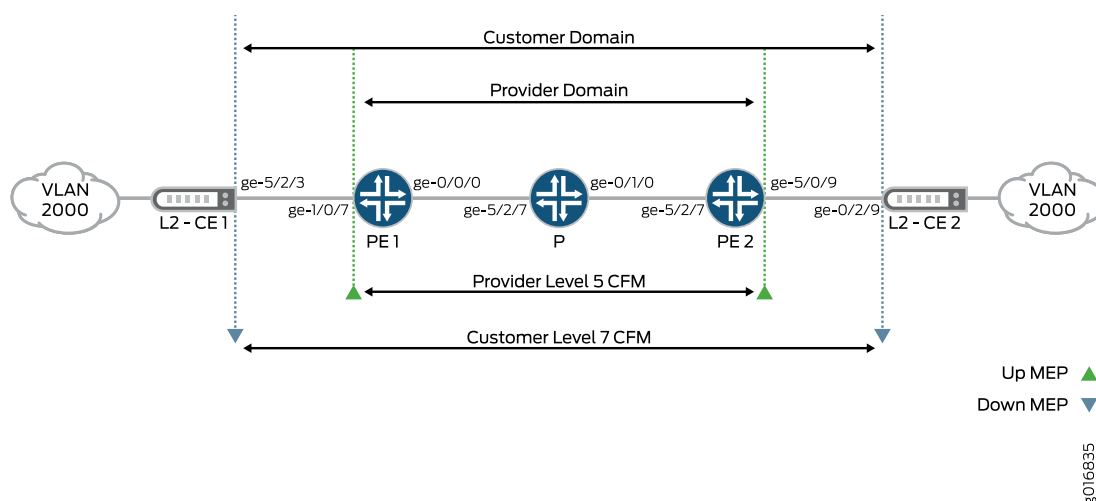
## Example: Configuring Ethernet CFM over VPLS

In this example, both the customer and service provider are running Ethernet CFM over a VPLS and a multiprotocol label switching (MPLS) network. The network is shown in [Figure 16 on page 614](#). The customer has configured Ethernet CFM on MX Series routers L2-CE1 and L2-CE2. The service provider has configured Ethernet CFM on MX Series routers PE1, P, and PE2.

**NOTE:** The configurations in this example are only partial examples of complete and functional router configurations. Do not copy these configurations and use them directly on an actual system.

The service provider is using CFM level 5 and the customer is using CFM level 7. The boundaries are marked with “up mep” and “down mep” CFM terminology in the figure.

Figure 16: Ethernet OAM with VPLS



**NOTE:** The logical interfaces in a VPLS routing instance might have the same or different VLAN configurations. VLAN normalization is required to switch packets correctly among these interfaces. Normalization supports automatic mapping of VLANs and performs operations on VLAN tags to achieve the desired translation. See *Configuring a Normalized VLAN for Translation or Tagging*.

**NOTE:**

The following forwarding path considerations must be observed:

- Packet receive path:
  - This is the forwarding path for packets received on the interfaces.
  - 802.1ag Ethernet OAM for VPLS uses implicit interface filters and forwarding table filters to flood, accept, and drop the CFM packets.
- Packet transmit path:
  - Junos OS uses the router's hardware-based forwarding for CPU-generated packets.
  - For down MEPs, the packets are transmitted on the interface on which the MEP is configured.
  - In MX series routers, for up MEPs, the packets must be flooded to other interfaces in the VPLS routing instance. The router creates a flood route tied to a flood next hop (with all interfaces to flood) and then sources the packets to be forwarded with this flood route.

The following are the configurations of the VPLS and CFM on the service provider routers.

**Configuration of PE1**

```
[edit chassis]
fpc 5 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}

[edit interfaces]
ge-1/0/7 {
  encapsulation flexible-ethernet-services;
  vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id 2000;
  }
}
ge-0/0/0 {
```

```

    unit 0 {
        family inet {
            address 10.200.1.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.168.231/32 {
                primary;
            }
            address 127.0.0.1/32;
        }
    }
}

[edit routing-instances]
vpls-vlan2000 {
    instance-type vpls;
    vlan-id 2000;
    interface ge-1/0/7.1;
    route-distinguisher 10.255.168.231:2000;
    vrf-target target:1000:1;
    protocols {
        vpls {
            site-range 10;
            site vlan2000-PE1 {
                site-identifier 2;
            }
        }
    }
}

[edit protocols]
rsvp {
    interface ge-0/0/0.0;
}
mpls {
    label-switched-path PE1-to-PE2 {
        to 10.100.1.1;
    }
}

```

```

    }
    interface ge-0/0/0.0;
  }
  bgp {
    group PE1-to-PE2 {
      type internal;
      local-address 10.200.1.1;
      family l2vpn {
        signaling;
      }
      local-as 65000;
      neighbor 10.100.1.1;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
      interface ge-0/0/0.0;
    }
  }
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain customer-site1 {
          level 5;
          maintenance-association customer-site1 {
            continuity-check {
              interval 1s;
            }
            mep 100 {
              interface ge-1/0/7.1;
              direction up;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}

```

```

    }
}

```

## Configuration of PE2

```

[edit chassis]
fpc 5 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}

[edit interfaces]
ge-5/0/9 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id 2000;
  }
}
ge-5/2/7 {
  unit 0 {
    family inet {
      address 10.100.1.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.168.230/32 {
        primary;
      }
      address 127.0.0.1/32;
    }
  }
}

```

```

}

[edit routing-instances]
vpls-vlan2000 {
    instance-type vpls;
    vlan-id 2000;
    interface ge-5/0/9.1;
    route-distinguisher 10.255.168.230:2000;
    vrf-target target:1000:1;
    protocols {
        vpls {
            site-range 10;
            site vlan2000-PE2 {
                site-identifier 1;
            }
        }
    }
}

[edit protocols]
rsvp {
    interface ge-5/2/7.0;
}
mpls {
    label-switched-path PE2-to-PE1 {
        to 10.200.1.1;
    }
    interface ge-5/2/7.0;
}
bgp {
    group PE2-to-PE1 {
        type internal;
        local-address 10.100.1.1;
        family l2vpn {
            signaling;
        }
        local-as 65000;
        neighbor 10.200.1.1;
    }
}
ospf {
    traffic-engineering;
}

```

```

reference-bandwidth 4g;
area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface ge-5/2/7.0;
}
}
oam {
    ethernet {
        connectivity-fault-management {
            maintenance-domain customer-site1 {
                level 5;
                maintenance-association customer-site1 {
                    continuity-check {
                        interval 1s;
                    }
                    mep 200 {
                        interface ge-5/0/9.1;
                        direction up;
                        auto-discovery;
                    }
                }
            }
        }
    }
}
}
}

```

### Configuration of P router

MPLS only, no CFM needed:

```

[edit]
interfaces {
    ge-5/2/7 {
        # Connected to PE1
        unit 0 {
            family inet {
                address 10.200.1.10/24;
            }
        }
    }
}

```



```

    }
    family mpls;
  }
}
ge-0/1/0 {
  # Connected to PE2
  unit 0 {
    family inet {
      address 10.100.1.10/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.168.240/32;
    }
  }
}
}

```

```

[edit]
protocols {
  rsvp {
    interface ge-0/1/0.0;
    interface ge-5/2/7.0;
  }
  mpls {
    interface ge-0/1/0.0;
    interface ge-5/2/7.0;
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
      interface ge-0/1/0.0;
      interface ge-5/2/7.0;
    }
  }
}

```

```

    }
  }
}

```

### CFM on L2-CE1

Here is the configuration of CFM on L2-E1:

```

[edit interfaces]
ge-5/2/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam]
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
        mep 800 {
          interface ge-5/2/3.0;
          direction down;
          auto-discovery;
        }
      }
    }
  }
}
}

```

### CFM on L2-CE2

Here is the configuration of CFM L2-CE2:

```
[edit interfaces]
ge-0/2/9 {
  vlan-tagging;
  unit 0 {
    vlan-id 2000;
  }
}

[edit protocols oam]
ethernet {
  connectivity-fault-management {
    maintenance-domain customer {
      level 7;
      maintenance-association customer-site1 {
        continuity-check {
          interval 1s;
        }
        mep 700 {
          interface ge-0/2/9.0;
          direction down;
          auto-discovery;
        }
      }
    }
  }
}
```

SEE ALSO

| [Configuring Continuity Check Messages](#) | 648

## Release History Table

| Release              | Description                                                                                                                                                                                                                                                                                              |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">17.1</a> | Starting in Release 17.1, Junos OS connectivity fault management (CFM), during a unified in-service software upgrade (ISSU), works when the peer device is not a Juniper Networks router.                                                                                                                |
| <a href="#">15.1</a> | Starting in Junos OS Release 15.1, configuring ISSU with CFM (802.1ag) is supported only on MX and PTX routers that support TLV.                                                                                                                                                                         |
| <a href="#">12.3</a> | Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 5G Universal Routing Platforms, you no longer need to configure the <b>no-control-word</b> statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. |
| <a href="#">12.3</a> | Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 5G Universal Routing Platforms, you no longer need to configure the <b>no-control-word</b> statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. |
| <a href="#">12.3</a> | Starting with Junos OS Release 12.3, for all interfaces configured on Modular Port Concentrators (MPCs) on MX Series 5G Universal Routing Platforms, you no longer need to configure the <b>no-control-word</b> statement for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs. |

## RELATED DOCUMENTATION

[Configuring Continuity Check Messages](#) | 648

[Inline Transmission Mode](#) | 824

## CFM Monitoring between CE and PE Devices

### IN THIS SECTION

- [Understanding CFM Monitoring between CE and PE Devices](#) | 625
- [Configuring Port Status TLV and Interface Status TLV](#) | 626
- [Configuring Chassis ID TLV](#) | 640

- [Configuring MAC Flush Message Processing in CET Mode | 642](#)
- [Example: Configuring an Action Profile Based on Connection Protection TLVs | 645](#)

Use this topic to understand more about CFM monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device. Also, you can understand more about how Interface Status TLVs, port status TLVs, chassis ID TLV, and connection protection TLV help in monitoring your network.

## Understanding CFM Monitoring between CE and PE Devices

You can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device. When the interface is down, CFM propagates the status of the interface in the CC messages. The CC message informs the customer edge device that the provider edge device is down.

You can configure CFM monitoring using either of the following two options:

- **Interface Status TLV (Type, Length, and Value)**—You can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device by using Interface Status TLV. When the interface is down, CFM propagates the status of the interface using interface status TLV. The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured, or the next-lower interface in the IETF RFC 2863 IF-MIB. Thus, the customer edge device is aware that the provider edge device is down. To configure CFM monitoring using Interface Status TLV, use the **interface-status-tlv** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain maintenance-domain maintenance-association maintenance-association continuity-check]** hierarchy level. This is the standard option.
- **RDI (Remote Defect Indication)**—Starting in Junos OS Release 17.3R1, you can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device by using the remote defect indication (RDI) bit. When you enable CFM monitoring, CFM propagates the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages. Thus, the customer edge device is aware that the provider edge device is down. The RDI bit is cleared when the service is back up. To configure CFM monitoring using the RDI bit, use the **interface-status-send-rdi** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain maintenance-domain maintenance-association maintenance-association continuity-check]** hierarchy level. This option is required if the customer edge device does not support Interface Status TLV.

**NOTE:** When the interface is set to CCC down and you have configured RDI, then RDI bit is sent. CFM does not monitor the status of the interface. If CCC down is set when the interface is not standby, RDI bit is sent with the CC messages if you have configured RDI.

### ***Single Active Multi-homing Use Case using RDI bit***

Consider the following topology where there are two provider edge devices (PE1 and PE2) as well as two customer edge devices (CE1 and CE2). PE1 is in active state while PE2 is in standby state. CFM down MEP is configured between the PE and CE. CFM detects that the CCC down and because CFM down MEP is configured, the CC messages generated have the RDI bit. The CC messages from PE2 to CE2 have the RDI bit set to indicate the blocked state. When PE2 becomes active, CCM down is cleared and the RDI bit is cleared from the subsequent CC messages.

### ***Active/Active Multihoming Use case using RDI bit***

Consider the topology where there are two provider edge devices (PE1 and PE2) and two customer edge devices (CE1 and CE2). PE1 is in active state while PE2 is in standby state. If CFM down MEP is not configured between the PE and CE to monitor the link connectivity, the CC messages generated do not have the RDI bit. CFM down MEP is configured between the PE and CE. CFM detects that the CCC down and because CFM down MEP is configured, the CC messages generated have the RDI bit. The CC messages from PE2 to CE2 have the RDI bit set to indicate the blocked state. When PE2 becomes active, CCM down is cleared and the RDI bit is cleared from the subsequent CC messages.

SEE ALSO

[interface-status-tlv | 1036](#)

[interface-status-send-rdi | 1037](#)

## **Configuring Port Status TLV and Interface Status TLV**

### **IN THIS SECTION**

- [TLVs Overview | 627](#)
- [Various TLVs for CFM PDUs | 627](#)
- [Support for Additional Optional TLVs | 629](#)
- [MAC Status Defects | 636](#)
- [Configuring Remote MEP Action Profile Support | 638](#)
- [Monitoring a Remote MEP Action Profile | 639](#)

### TLVs Overview

Type, Length, and Value (TLVs) are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and/or optional information in a PDU. TLVs are not aligned to any particular word or octet boundary. TLVs follow each other with no padding between them.

[Table 85 on page 627](#) shows the TLV format and indicates if it is required or optional.

**Table 85: Format of TLVs**

| Parameter | Octet (sequence) | Description                                                                                                                                                                                                              |
|-----------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type      | 1                | Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.                                                                                                            |
| Length    | 2–3              | Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field. |
| Value     | 4                | Length specified by the Length field. Optional. Not present if the Type field is 0 or if the Length field is 0.                                                                                                          |

### Various TLVs for CFM PDUs

[Table 86 on page 627](#) shows a set of TLVs defined by IEEE 802.1ag for various CFM PDU types. Each TLV can be identified by the unique value assigned to its type field. Some type field values are reserved.

**Table 86: Type Field Values for Various TLVs for CFM PDUs**

| TLV or Organization       | Type Field |
|---------------------------|------------|
| End TLV                   | 0          |
| Sender ID TLV             | 1          |
| Port Status TLV           | 2          |
| Data TLV                  | 3          |
| Interface Status TLV      | 4          |
| Reply Ingress TLV         | 5          |
| Reply Egress TLV          | 6          |
| LTM Egress Identifier TLV | 7          |
| LTR Egress Identifier TLV | 8          |

Table 86: Type Field Values for Various TLVs for CFM PDUs (*continued*)

| TLV or Organization       | Type Field |
|---------------------------|------------|
| Reserved for IEEE 802.1   | 9 to 30    |
| Organization-Specific TLV | 31         |
| Defined by ITU-T Y.1731   | 32 to 63   |
| Reserved for IEEE 802.1   | 64 to 255  |

Not every TLV is applicable for all types of CFM PDUs.

- TLVs applicable for continuity check message (CCM):

- End TLV
- Sender ID TLV
- Port Status TLV
- Interface Status TLV
- Organization-Specific TLV

- TLVs applicable for loopback message (LBM):

- End TLV
- Sender ID TLV
- Data TLV
- Organization-Specific TLV

- TLVs applicable for loopback reply (LBR):

- End TLV
- Sender ID TLV
- Data TLV
- Organization-Specific TLV

- TLVs applicable for linktrace message (LTM):

- End TLV
- LTM Egress Identifier TLV



- Sender ID TLV
- Organization-Specific TLV
- TLVs applicable for linktrace reply (LTR):
  - End TLV
  - LTR Egress Identifier TLV
  - Reply Ingress TLV
  - Reply Egress TLV
  - Sender ID TLV
  - Organization-Specific TLV

The following TLVs are currently supported in the applicable CFM PDUs:

- End TLV
- Reply Ingress TLV
- Reply Egress TLV
- LTR Egress Identifier TLV
- LTM Egress Identifier TLV
- Data TLV

#### ***Support for Additional Optional TLVs***

##### **IN THIS SECTION**

- [Port Status TLV | 630](#)
- [Interface Status TLV | 633](#)

The following additional optional TLVs are supported:

- Port Status TLV
- Interface Status TLV

MX Series routers support configuration of port status TLV and interface status TLV. Configuring the Port Status TLV allows the operator to control the transmission of the Port Status TLV in CFM PDUs.

**NOTE:** Although Port Status TLV configuration statements are visible in the CLI on M120 and M320 routers, Port Status TLV cannot be configured on these systems. Port Status TLV can be enabled on a MEP interface only if it is a bridge logical interface, which is not possible on these systems.

For configuration information, see the following sections:

### **Port Status TLV**

The Port Status TLV indicates the ability of the bridge port on which the transmitting MEP resides to pass ordinary data, regardless of the status of the MAC. The value of this TLV is driven by the MEP variable **enableRmepDefect**, as shown in [Table 88 on page 630](#). The format of this TLV is shown in [Table 87 on page 630](#).

Any change in the Port Status TLVs value triggers one extra transmission of that bridge ports MEP CCMs.

**Table 87: Port Status TLV Format**

| Parameter                                         | Octet (Sequence) |
|---------------------------------------------------|------------------|
| Type = 2                                          | 1                |
| Length                                            | 2-3              |
| Value (See <a href="#">Table 88 on page 630</a> ) | 4                |

**Table 88: Port Status TLV Values**

| Mnemonic  | Ordinary Data Passing Freely Through the Port | Value |
|-----------|-----------------------------------------------|-------|
| psBlocked | No: <b>enableRmepDefect</b> = false           | 1     |
| psUp      | Yes: <b>enableRmepDefect</b> = true           | 2     |

The MEP variable **enableRmepDefect** is a boolean variable indicating whether frames on the service instance monitored by the maintenance associations if this MEP are enabled to pass through this bridge port by the Spanning Tree Protocol and VLAN topology management. It is set to TRUE if:

- The bridge port is set in a state where the traffic can pass through it.
- The bridge port is running multiple instances of the spanning tree.
- The MEP interface is not associated with a bridging domain.

### Configuring Port Status TLV

Junos OS provides configuration support for the Port Status TLV, allowing you to control the transmission of this TLV in CCM PDUs. The Junos OS provides this configuration at the continuity-check level. By default, the CCM does not include the Port Status TLV. To configure the Port Status TLV, use the **port-status-tlv** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check]** hierarchy level.

**NOTE:** Port Status TLV configuration is not mandated by IEEE 802.1ag. The Junos OS provides it in order to give more flexibility to the operator; however it receives and processes CCMs with a Port Status TLV, regardless of this configuration.

An example of the configuration statements follows:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
          level number;
          maintenance-association identifier {
            continuity-check {
              interval number,
              loss-threshold number;
              hold-interval number;
              port-status-tlv; # Sets Port Status TLV
            }
          }
        }
      }
    }
  }
}
```

You cannot enable Port Status TLV transmission in the following two cases:

- If the MEP interface under the maintenance-association is not of type bridge.
- If the MEP is configured on a physical interface.

### Displaying the Received Port Status TLV

The Junos OS saves the last received Port Status TLV from a remote MEP. If the received Port Status value does not correspond to one of the standard values listed in [Table 88 on page 630](#), then the **show** command displays it as "unknown." You can display the last saved received Port Status TLV using the **show oam**

ethernet connectivity-fault-management mep-database maintenance-domain *identifier*

maintenance-association *identifier* local-mep *identifier* remote-mep *identifier* command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
```

```
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none # RX PORT STATUS
Interface status TLV: none
```

### ***Displaying the Transmitted Port Status TLV***

The Junos OS saves the last transmitted Port Status TLV from a local MEP. If the transmission of the Port Status TLV has not been enabled, then the **show** command displays "none." You can display the last saved transmitted Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
```

```
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up # TX PORT STATUS
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
```

```
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
```

**Interface Status TLV**

The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured, or the next-lower interface in the IETF RFC 2863 IF-MIB. The format of this TLV is shown in [Table 89 on page 633](#). The enumerated values are shown in [Table 90 on page 633](#).

**Table 89: Interface Status TLV Format**

| Parameter                                         | Octet (Sequence) |
|---------------------------------------------------|------------------|
| Type = 4                                          | 1                |
| Length                                            | 2-3              |
| Value (See <a href="#">Table 90 on page 633</a> ) | 4                |

**Table 90: Interface Status TLV Values**

| Mnemonic         | Interface Status | Value |
|------------------|------------------|-------|
| isUp             | up               | 1     |
| isDown           | down             | 2     |
| isTesting        | testing          | 3     |
| isUnknown        | unknown          | 4     |
| isDormant        | dormant          | 5     |
| isNotPresent     | notPresent       | 6     |
| isLowerLayerDown | lowerLayerDown   | 7     |

**NOTE:** When the operational status of a logical interface changes from the down state (status value of 2) to the lower layer down state (status value of 7) and vice versa, the LinkDown SNMP trap is not generated. For example, if you configure an aggregated Ethernet interface bundle with a VLAN tag and add a physical interface that is in the operationally down state to the bundle, the operational status of the aggregated Ethernet logical interface bundle at that point is lower layer down (7). If you take the MIC associated with the interface offline, the LinkDown trap is not generated when the logical interface shifts from the lower layer down state to the down state.

Similarly, consider another sample scenario in which an physical interface is added to an aggregated Ethernet bundle that has VLAN tagging and the aggregated Ethernet logical interface is disabled. When the logical interface is disabled, the operational status of the logical interface changes to down. If you disable the physical interface that is part of the aggregated Ethernet bundle, the operational status of the aggregated Ethernet logical interface remains down. If you reenables the aggregated Ethernet logical interface, the operational status of it changes from down to lower layer down. The LinkDown SNMP trap is not generated at this point.

### Configuring Interface Status TLV

The Junos OS provides configuration support for the Interface Status TLV, thereby allowing operators to control the transmission of this TLV in CCM PDUs through configuration at the continuity-check level.

**NOTE:** This configuration is not mandated by IEEE 802.1ag; rather it is provided to give more flexibility to the operator. The Junos OS receives and processes CCMs with the Interface Status TLV, regardless of this configuration.

The interface status TLV configuration is shown below:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
          level number;
          maintenance-association identifier {
            continuity-check {
              interval number;
              loss-threshold number;
              hold-interval number;
            }
          }
        }
      }
    }
  }
}
```

```

        interface-status-tlv; # Sets the interface status TLV
    }
}
}
}
}
}
}
}

```

**NOTE:** The Junos OS supports transmission of only three out of seven possible values for the Interface Status TLV. The supported values are 1, 2, and 7. However, the Junos OS is capable of receiving any value for the Interface Status TLV.

### Displaying the Received Interface Status TLV

The Junos OS saves the last received Interface Status TLV from the remote MEP. If the received Interface Status value does not correspond to one of the standard values listed in [Table 89 on page 633](#), then the **show** command displays "unknown."

You can display this last saved Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```

user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md5 maintenance-association ma5 local-mep 2001 remote-mep 1001

```

```

Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false

```

```
Port status TLV: none
Interface status TLV: none # displays the Interface Status TLV state
```

### **Displaying the Transmitted Interface Status TLV**

The Junos OS saves the last transmitted Interface Status TLV from a local MEP. If the transmission of Interface Status TLV has not been enabled, then the **show** command displays "none."

You can display the last transmitted Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
```

```
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
```

### **MAC Status Defects**

The Junos OS provides MAC status defect information, indicating that one or more of the remote MEPs is reporting a failure in its Port Status TLV or Interface Status TLV. It indicates "yes" if either some remote MEP is reporting that its interface is not isUp (for example, at least one remote MEPs interface is unavailable), or if all remote MEPs are reporting a Port Status TLV that contains some value other than psUp (for example, all remote MEPs Bridge Ports are not forwarding data). There are two **show** commands you can use to view the MAC Status Defects indication.

Use the **mep-database** command to display MAC status defects:



```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md6 maintenance-association ma6
```

```
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM          : no
  Erroneous CCM received                 : no
  Cross-connect CCM received             : no
  RDI sent by some MEP                  : no
  Some remote MEP's MAC in error state   : yes # MAC Status Defects yes/no
Statistics:
  CCMS sent                             : 1658
  CCMS received out of sequence          : 0
  LBMS sent                             : 0
  Valid in-order LBRs received           : 0
  Valid out-of-order LBRs received       : 0
  LBRs received with corrupted data      : 0
  LBRs sent                             : 0
  LTMS sent                             : 0
  LTMS received                         : 0
  LTRs sent                             : 0
  LTRs received                         : 0
  Sequence number of next LTM request    : 0
  1DMs sent                             : 0
  Valid 1DMs received                   : 0
  Invalid 1DMs received                  : 0
  DMMs sent                             : 0
  DMRs sent                             : 0
  Valid DMRs received                   : 0
  Invalid DMRs received                  : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    200      00:05:85:73:39:4a  ok    xe-5/0/0.0
```

Use the **interfaces** command to display MAC status defects:

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
```

```

Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
Interface status TLV: up, Port status TLV: up
MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
MEP status: running
Defects:
    Remote MEP not receiving CCM                : no
    Erroneous CCM received                       : no
    Cross-connect CCM received                   : no
    RDI sent by some MEP                        : no
    Some remote MEP's MAC in error state         : yes # MAC Status Defects
yes/no
Statistics:
    CCMs sent                                   : 1328
    CCMs received out of sequence                : 0
    LBMs sent                                   : 0
    Valid in-order LBRs received                 : 0
    Valid out-of-order LBRs received             : 0
    LBRs received with corrupted data            : 0
    LBRs sent                                   : 0
    LTMs sent                                   : 0
    LTMs received                               : 0
    LTRs sent                                   : 0
    LTRs received                               : 0
    Sequence number of next LTM request          : 0
    1DMs sent                                   : 0
    Valid 1DMs received                         : 0
    Invalid 1DMs received                       : 0
    DMMs sent                                   : 0
    DMRs sent                                   : 0
    Valid DMRs received                         : 0
    Invalid DMRs received                       : 0
Remote MEP count: 1
    Identifier  MAC address  State  Interface
    200        00:05:85:73:39:4a  ok    xe-5/0/0.0

```

### Configuring Remote MEP Action Profile Support

Based on values of **interface-status-tlv** and **port-status-tlv** in the received CCM packets, a specific action, such as **interface-down**, can be taken using the **action-profile** options. Multiple action profiles can be configured on the router, but only one action profile can be assigned to a remote MEP.

The action profile can be configured with at least one event to trigger the action; but the action will be triggered if any one of these events occurs. It is not necessary for all of the configured events to occur to trigger **action**.

An action-profile can be applied only at the remote MEP level.

The following example shows an action profile configuration with explanatory comments added:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile tlv-action {
  event {
    # If interface status tlv with value specified in the config is received
    interface-status-tlv down|lower-layer-down;
    # If port status tlv with value specified in the config is received
    port-status-tlv blocked;
    # If connectivity is lost to the peer */
    adjacency-loss;
  }
  action {
    # Bring the interface down */
    interface-down;
  }
  default-actions interface-down;
}
# domains
maintenance-domain identifier {
  # maintenance domain level (0-7)
  level number;
  # association
  maintenance-association identifier {
    mep identifier {
      interface ge-x/y/z.w;
      remote-mep identifier {
        # Apply the action-profile for the remote MEP
        action-profile tlv-action;
      }
    }
  }
}
```

### Monitoring a Remote MEP Action Profile

You can use the **show oam ethernet connectivity-fault-management mep-database** command to view the action profile status of a remote MEP, as in the following example:

```
show oam ethernet connectivity-fault-  
management mep-database remote-mep
```

**(Action Profile Event)**

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md5 maintenance-association ma5 remote-mep 200
```

```
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none # last status TLVs transmitted
by the router
Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

Remote MEP identifier: 200, State: ok # displays the remote MEP name and state
MAC address: 00:05:85:73:96:1f, Type: Configured
Interface: ge-1/0/8.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: lower-layer-down
Action profile: juniper # displays remote MEP's action profile identifier
Last event: Interface-status-tlv lower-layer-down # last remote MEP event

# to trigger action
Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
# action occurrence time
```

**RELATED DOCUMENTATION**

[connectivity-fault-management | 1132](#)

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)

**Configuring Chassis ID TLV**

In Release 16.1R2 and later, you can configure Junos OS to send the sender ID TLV along with the packets. The sender ID TLV is an optional TLV that is sent in continuity check messages (CCMs), loopback messages, and Link Trace Messages (LTMs), as specified in the IEEE 802.1ag standard. The sender ID TLV contains the chassis ID, which is the unique, CFM-based MAC address of the device, and the management IP address, which is an IPv4 or an IPv6 address.

The value of the **length** field in the TLV indicates whether or not the TLV contains the chassis ID information. The possible values for the **length** field are zero (0) or any valid number, which indicates the absence or presence of chassis ID information in the TLV, respectively.

You can enable Junos OS to send the sender ID TLV at the global level by using the **set protocols oam ethernet connectivity-fault-management sendid-tlv send-chassis-tlv** command. If the sender ID TLV is configured at the global level, then the default maintenance domain, maintenance association, and the maintenance association intermediate point (MIP) half function inherit this configuration.

You can also configure the sender ID TLV at the following hierarchy levels:

- **[edit protocols oam ethernet connectivity-fault-management]**.
- **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name* maintenance-association *maintenance-association-name* continuity-check]**.

The sender ID TLV configuration at the maintenance-association level takes precedence over the global-level configuration.

**NOTE:** The sender ID TLV is supported only for 802.1ag PDUs and is not supported for performance monitoring protocol data units (PDUs).

SEE ALSO

| [IEEE 802.1ag OAM Connectivity Fault Management Overview](#) | 548

## Configuring MAC Flush Message Processing in CET Mode

In carrier Ethernet transport (CET) mode, MX Series routers are used as provider edge (PE) routers, and Nokia Siemens Networks A2200 Carrier Ethernet Switches (referred to as E-domain devices) that run standard-based protocols are used in the access side. On the MX Series routers, VPLS pseudowires are configured dynamically through label distribution protocol (LDP). On the E-domain devices, topology changes are detected through connectivity fault management (CFM) sessions running between the E-domain devices and the MX Series PE routers. The MX Series PE routers can bring the carrier Ethernet interface down if there is CFM connectivity loss. This triggers a local MAC flush as well as a targeted label distribution protocol (T-LDP) MAC flush notification that gets sent towards the remote MX Series PEs to trigger MAC flush on them.

In CET inter-op mode, MX Series routers need to interoperate with the Nokia Siemens Networks Ax100 Carrier Ethernet access devices (referred to as A-domain devices) that run legacy protocols. Nokia Siemens Networks A4100 and A8100 devices act as an intermediate between the MX Series PE routers and A-domain devices. These intermediate devices perform interworking function (IWF) procedures so that operations administration management (OAM) sessions can be run between MX Series routers and A-domain devices. There are no VPLS pseudowires between the MX Series PE routers and the Nokia Siemens Networks A4100 and A8100 intermediate devices, so there is no LDP protocol running between the PE routers to send topology change notifications. In order to communicate topology changes, MX Series routers can trigger a MAC flush and propagate it in the core. MX Series routers can use action profiles based upon the connection protection type length value (TLV) event. The action profile brings down the carrier edge logical interface in MX Series PE routers, which will trigger a local MAC flush and also propagate the topology change to the core using LDP notification.

For VPLS there is no end-to-end connectivity monitored. The access rings are independently monitored by running CFM down multiple end points (MEPs) on the working and protection paths for each of the services between the E-domain devices and the MX Series PE routers, and between the A-domain devices and the MX Series PE routers the IWF hosted by the Nokia Siemens Networks A-4100 devices. When there is a connectivity failure on the working path, the Nokia Siemens Networks Ax200 devices perform a switchover to the protection path, triggering a topology change notification (in the form of TLVs carried in CCM) to be sent on the active path.

Figure 17: CET inter-op Dual Homed Topology




Figure 17 on page 643 describes the dual homed topology on MX Series PE routers connected to the A-domain. When an A-domain device triggers a switchover, it starts switching the service traffic to the new active path. This change is communicated in the HELLO protocol data units (PDUs) sent by that A-domain device on the working and protection paths. When the IWF in A4100 receives these HELLO PDUs, it converts them to standard CCM messages and also inserts a connection protection TLV. The “Protection-in-use” field of the connection protection TLV is encoded with the currently active path, and is included in the CCM message. CCM messages are received by the MX Series PE routers through the VLAN spoke in A4100. In the above dual homed scenario, one MX Series PE router monitors the working path, and the other MX Series PE router monitors the protection path.

A MAC flush occurs when the CFM session that is monitoring the working path detects that the service traffic has moved to the protection path or when the CFM session that is monitoring the protection path detects that the service traffic has moved to the working path.

Figure 18: CET inter-op Dual Attached Topology

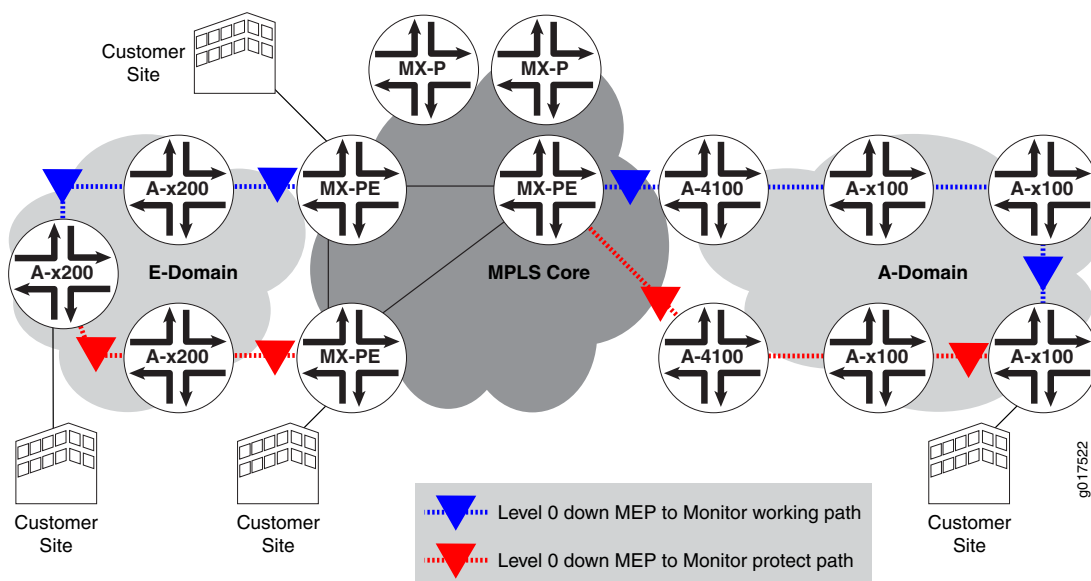


Figure 18 on page 644 describes the dual attached topology on MX Series PE routers connected to the A-domain. The MAC flush mechanism used in this case is also the same as the one used for the A-domain in the dual homed scenario (Figure 1). However in this case both the CFM sessions are hosted by only one MX Series PE router. When Ax100 in the A-domain detects topology changes, the MX Series PE router receives the connection protection TLV in the CCM message for the working and protection paths with the value of “Protection-in-use” indicating which path is the active one. Based upon the event that is generated for the CFM session, the MX Series PE router will bring down the appropriate interface which will trigger a local MAC flush.

### Configuring a Connection Protection TLV Action Profile

An action profile can be configured to perform the **interface-down** action based on the values of **connection-protection-tlv** in the received CCM packets.

The following example shows an action profile configuration with explanatory comments added:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
    # If a connection protection TLV with a "Protection-in-use" value of SET is received */
    connection-protection-tlv <using-protection-path>;
    # If a connection protection TLV with a "Protection-in-use" value of RESET is received */
    connection-protection-tlv <using-working-path>;
  }
  action {
    # Bring the interface down */
    interface-down;
```



```
}
}
```

## SEE ALSO

[connection-protection-tlv | 1010](#)

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)

## Example: Configuring an Action Profile Based on Connection Protection TLVs

### IN THIS SECTION

- [Requirements | 645](#)
- [Overview and Topology | 645](#)
- [Configuration | 646](#)

This example shows how to configure an action profile based on the connection protection TLV for the purposes of triggering MAC flushes based on topology changes in a CET network.

### **Requirements**


This example uses the following hardware and software components:

- Junos OS Release 11.2 or later
- A MX series PE router

### **Overview and Topology**

The physical topology of a CET network using MX series PE routers is shown in [Figure 19 on page 646](#)

Figure 19: Topology of CET network



The following definitions describe the meaning of the device abbreviation and terms used in [Figure 19 on page 646](#).

- Provider edge (PE) device—A device, or set of devices, at the edge of the provider network that presents the provider's view of the customer site.
- E-domain—Nokia Siemens Networks Carrier Ethernet Switches that run standard based protocols and are used in the access side.
- A-domain—Nokia Siemens Networks Carrier Ethernet Switches that run legacy protocols.

### Configuration

#### Step-by-Step Procedure

To configure an action profile based on the connection protection TLV, perform these tasks:

1. Configure an action profile

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile <tlv-action> {
  event {
```

2. If the connection protection TLV is received with a "Protection-in-use" value of SET, then the connection protection TLV should use the protection path

```
connection-protection-tlv <using-protection-path>;
```

3. If the connection protection TLV is received with a “Protection-in-use” value of RESET, then the connection protection TLV should use the working path

```
connection-protection-tlv <using-working-path>;  
}
```

4. Configure the action profile to bring the interface down

```
action {  
    /* Bring the interface down */  
    interface-down;  
}  
}
```

## Results

Check the results of the configuration

```
[edit protocols oam ethernet connectivity-fault-management]  
action-profile <tlv-action> {  
    event {  
        connection-protection-tlv <using-protection-path>;  
        connection-protection-tlv <using-working-path>;  
    }  
    action {  
        interface-down;  
    }  
}
```

## SEE ALSO

| [connection-protection-tlv](#) | 1010

Release History Table

| Release                | Description                                                                                                                                                                                                                                                       |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">17.3R1</a> | Starting in Junos OS Release 17.3R1, you can enable connectivity fault management (CFM) monitoring between provider edge devices and customer edge devices when the customer edge device is not a Juniper device by using the remote defect indication (RDI) bit. |
| <a href="#">16.1</a>   | In Release 16.1R2 and later, you can configure Junos OS to send the sender ID TLV along with the packets.                                                                                                                                                         |

RELATED DOCUMENTATION

|                                                                               |
|-------------------------------------------------------------------------------|
| <a href="#">Introduction to OAM Connectivity Fault Management (CFM)   546</a> |
| <a href="#">ITU-T Y.1731 Ethernet Service OAM Overview   689</a>              |

## Configuring Continuity Check Messages

IN THIS SECTION

- [Configuring Faster Protection Switching for Point-to-Point Network Topologies | 649](#)
- [Configuring Faster Convergence for Dual-Homed Multipoint-to-Multipoint Network Topologies | 651](#)
- [Configuring a Primary VLAN ID for Increased Flexibility | 652](#)
- [Configuring a Remote Maintenance Association to Accept a Different ID | 653](#)

Junos OS provides enhancements to trigger faster protection-switching and convergence in the event of failures in Ethernet domains for Carrier Ethernet services. These enhancements can be used when CE devices in the Ethernet domain detect faster service failures and propagates the information in the interface-status TLV of the continuity-check messages (CCMs). When CCMs are received, PE devices can perform certain actions which facilitates faster protection-switching and convergence. You can configure CCM for better scalability using the information provided in this topic.

## Configuring Faster Protection Switching for Point-to-Point Network Topologies

You can apply an action profile to provide faster protection switching for point-to-point network topologies with local switching configured. In a normal state, CCM sessions are configured on the working and protect interfaces. The CCM packets transmitted contain an interface-status TLV with the value up on the working interface and value down on the protect interface. When a link fails on the working interface, the protect interface starts receiving the interface-status TLV as up. With the profile configuration, if the interface-status TLV received on the protect interface is up, the working interface is automatically marked as **interface-down**.

To configure the **interface-status-tlv** down event, include the **interface-status-tlv down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure **interface-down** as the action profile's action, include the **interface-down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* action]** hierarchy level.

To configure **interface-down *peer-interface*** as the clear-action, include **interface-down *peer-interface*** at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* clear-action]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile p1 {
      event {
        interface-status-tlv down;
      }
      action {
        interface-down;
      }
      clear-action {
        interface-down peer-interface;
      }
    }
  }
}
```

In this action profile configuration, when the interface-status TLV is received as up, the *peer-interface* is marked as down.

The *peer-interface* is configured in the **protect-maintenance-association** statement. Consider the following example using the **protect-maintenance-association** statement in the configuration:

```

[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile p1 {
      event {
        adjacency-loss;
      }
      action {
        interface-down;
      }
      clear-action {
        interface-down peer-interface;
      }
    }
  }
  maintenance-domain nsn {
    level 5;
    maintenance-association ma1 {
      protect-maintenance-association ma2;
      continuity-check {
        interval 100ms;
        connection-protection-tlv;
      }
      mep 100 {
        interface ge-1/1/0.0;
        direction down;
        auto-discovery;
      }
    }
    maintenance-association ma2 {
      continuity-check {
        interval 100ms;
        connection-protection-tlv;
      }
    }
    mep 101 {
      interface ge-1/2/0.0;
      direction down;
      auto-discovery;
    }
    remote-mep 100
      action-profile p1;
  }
}
}

```

SEE ALSO

| [connectivity-fault-management](#) | 1132

## Configuring Faster Convergence for Dual-Homed Multipoint-to-Multipoint Network Topologies

You can apply an action profile to provide faster convergence for dual-homed multipoint-to-multipoint network topologies. If a multipoint-to-multipoint Ethernet service uses MAC-based forwarding and stale MAC addresses exist in the learning tables, this can result in traffic black holes in the network where incoming traffic is silently discarded, without informing the source that the data did not reach its intended recipient. With the profile configuration, if the interface-status TLV received on the protect interface is up, then the interface-status TLV on the working interface is marked as down and the PE device for the protect interface propagates a remote MAC-flush message to the PE devices in the virtual private LAN service (VPLS) by using TLDP-MAC-FLUSH. The MAC flush avoids traffic blackholing due to stale mac-db entries.

To configure the **interface-status-tlv** down event, include the **interface-status-tlv down** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* event]** hierarchy level.

To configure **propagate-remote-flush** as the action profile's action, include the **propagate-remote-flush** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* action]** hierarchy level.

To configure **propagate-remote-flush** as the clear-action, include the **propagate-remote-flush** statement at the **[edit protocols oam ethernet connectivity-fault-management action-profile *profile-name* clear-action]** hierarchy level.

```
[edit protocols oam]
ethernet {
  connectivity-fault-management {
    action-profile test {
      event {
        interface-status-tlv down;
      }
      action {
        propagate-remote-mac-flush;
      }
      clear-action {
        propagate-remote-mac-flush;
      }
    }
  }
}
```

In this action profile configuration, when the incoming CCM packet contains the interface-status TLV with value down, the **propagate-remote-mac-flush** action is triggered for the action-profile.

SEE ALSO

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)

[connectivity-fault-management | 1132](#)

## Configuring a Primary VLAN ID for Increased Flexibility

You can assign a primary virtual LAN (VLAN) ID in the maintenance association for increased flexibility in the number of tags. When a **vlan-range** or **vlan-id-list** is configured on an interface, the service OAM must run on one of the VLANs. The VLAN assigned for service monitoring is considered the primary VLAN. If a **primary-vid** is not configured, Junos OS assigns the first VLAN from the **vlan-range** or **vlan-id-list**. In earlier releases, Junos OS assigned VLAN 4095.

To configure a primary VLAN ID, you can specify the **primary-vid** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
maintenance domain md3 {
  level 3;
  maintenance-association ma3 {
    primary-vid 2000;
    continuity-check {
      interval 10ms;
      connection-protection-tlv;
    }
    mep 2 {
      interface ge-2/2/0.0;
      direction up;
      auto-discovery;
    }
  }
}
```

SEE ALSO

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)



---

[connection-protection-tlv | 1010](#)


---

[connectivity-fault-management | 1132](#)

## Configuring a Remote Maintenance Association to Accept a Different ID

You can configure a maintenance association to accept a different maintenance association identifier (ID) from a neighbor by including a **remote-maintenance-association** statement. The 802.1ag CCM sessions expect the same maintenance association identifier from its neighbors. If there is a maintenance association identifier mismatch, the PDUs are marked as error PDUs. If a **remote-maintenance-association** statement is configured, a different maintenance association identifier is accepted and the 802.1ag CCM sessions do not mark the CCM PDUs as error PDUs when the maintenance-association name is the same as the name specified in the **remote-maintenance-association** statement.

To configure a remote maintenance association, include the **remote-maintenance-association** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name*]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
maintenance domain md3 {
  level 1;
  maintenance-association ma3 {
    remote-maintenance-association fix-ma;
    continuity-check {
      interval 10ms;
      connection-protection-tlv;
    }
    mep 2 {
      interface ge-2/2/0.0;
      direction up;
      auto-discovery;
    }
  }
}
```

Using this configuration, interoperability is improved for CCMs with low-end CE devices supporting fixed maintenance association identifier configurations.

### SEE ALSO

---

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)


---

[connectivity-fault-management | 1132](#)


---

[connection-protection-tlv | 1010](#)

## RELATED DOCUMENTATION

[Introduction to OAM Connectivity Fault Management \(CFM\) | 546](#)

[Configuring Connectivity Fault Management \(CFM\) | 553](#)

# Configuring Link Fault Management

## IN THIS CHAPTER

- [Introduction to OAM Link Fault Management \(LFM\) | 655](#)
- [Configuring Link Fault Management | 660](#)
- [Remote Fault Detection for Link Fault Management | 678](#)
- [Remote Loopback for Link Fault Management | 680](#)

## Introduction to OAM Link Fault Management (LFM)

### IN THIS SECTION

- [IEEE 802.3ah OAM Link-Fault Management Overview | 655](#)
- [Understanding Ethernet OAM Link Fault Management for ACX Series Routers | 656](#)
- [Configuring Ethernet 802.3ah OAM | 658](#)

Use this topic to understand more about link fault management (LFM) and how to configure link fault management.

### IEEE 802.3ah OAM Link-Fault Management Overview

Ethernet interfaces capable of running at 100 Mbps or faster on EX Series switches, PTX Series, MX Series, M Series (except M5 and M10 routers), and T Series routers support the IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to being a WAN and access technology, as well as being backward-compatible with existing Ethernet technology. Junos OS supports IEEE 802.3ah link-fault management.

The features of link-fault management are:

- Discovery
- Link monitoring
- Remote fault detection
- Remote loopback

Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured.

The following features are not supported:

- Ethernet running on top of a Layer 2 protocol, such as Ethernet over ATM, is not supported in OAM configurations.
- Remote loopback is not supported on the 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The remote loopback feature mentioned in section 57.2.11 of IEEE 802.3ah is not supported on T4000 routers.

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

SEE ALSO

[Detecting Remote Faults | 678](#)

[Enabling Nonstop Routing for Ethernet Link Fault Management on Backup Routers | 682](#)

## Understanding Ethernet OAM Link Fault Management for ACX Series Routers

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

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

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

The following OAM LFM features are supported on ACX Series routers:

- **Discovery and Link Monitoring**

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in the process. The router performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

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

- **Remote Fault Detection**

Remote fault detection uses flags and events. Flags are used to convey the following:

- **Link Fault** means a loss of signal
- **Dying Gasp** means an unrecoverable condition such as a power failure. In this condition, the local peer informs the remote peer about the failure state. When the remote peer receives a dying-gasp PDU, it takes an action corresponding to the action profile configured with the **link-adjacency-loss** event.

**NOTE:** ACX5096 and ACX5048 routers do not support dying-gasp.

ACX Series routers can generate and receive dying-gasp packets. When LFM is configured on an interface, a dying-gasp PDU is generated for the interface on the following failure conditions:

- Power failure
- Packet Forwarding Engine panic or a crash
- **Critical Event** means an unspecified vendor-specific critical event.

You can specify the interval at which OAM PDUs are sent for fault detection.

**NOTE:** ACX Series routers support the receipt of dying-gasp packets, but cannot generate them.

- Remote Loopback Mode

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

If a remote data terminal equipment (DTE) supports remote loopback mode, Junos OS can place the remote DTE into loopback mode. When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs and PAUSE frames are looped back. No changes are made to the frames. OAM PDUs continue to be sent and processed.

SEE ALSO

| [Remote Fault Detection for Link Fault Management](#) | 678

## Configuring Ethernet 802.3ah OAM

The IEEE 802.3ah standard for Operation, Administration, and Management (OAM) provides a specification for *Ethernet in the first mile (EFM)* connectivity. EFM defines how Ethernet can be transmitted over new media types using new Ethernet physical layer (PHY) interfaces. You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah OAM standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to being a WAN and access technology, as well as being backward-compatible with existing Ethernet technology.

For Ethernet interfaces capable of running at 100 Mbps or faster, the IEEE 802.3ah OAM standard is supported on numerous Juniper Networks routers and switches. This topic describes configuration support for IEEE 802.3ah OAM features on routers.

Beginning in Junos OS Release 12.1, PTX Series routers support the following IEEE 802.3ah OAM features at the physical interface level:

- Discovery and link monitoring
- Fault signaling and detection
- Periodic packet management (PPM) processing

- Action profile support
- Graceful Routing Engine switchover (GRES)

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

```
oam {
  ethernet {
    link-fault-management {
      interfaces {
        interface-name {
          pdu-interval interval;
          link-discovery (active | passive);
          pdu-threshold count;
        }
      }
    }
  }
}
```

You can configure threshold values for fault events that trigger the sending of link event TLVs when the values exceed the threshold. To set threshold values for fault events on an interface, include the **event-thresholds** statement at the **[edit protocols oam ethernet link-fault-management interface]** hierarchy level.

You can also configure OAM threshold values within an action profile and apply the action profile to multiple interfaces. To create an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level.

You can configure Ethernet OAM either on an aggregate interface or on each of its member links. However, we recommend that you configure Ethernet OAM on the aggregate interface, and this will internally enable Ethernet OAM on the member links.

To view OAM statistics, use the **show oam ethernet link-fault-management** operational mode command. To clear OAM statistics, use the **clear oam ethernet link-fault-management statistics** operational mode command. To clear link-fault management state information and restart the link discovery process on Ethernet interfaces, use the **clear oam ethernet link-fault-management state** operational mode command. For more information about these commands, see the [CLI Explorer](#).

To enable IEEE 802.3ah OAM support, include the **interface** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

When you enable IEEE 802.3ah OAM on a physical interface, the discovery process is automatically triggered.

SEE ALSO

[event-thresholds](#) | [1175](#)

[action-profile](#) | [994](#)

Release History Table

| Release                | Description                                                                                                                                                                                   |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">17.3R1</a> | Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured. |

RELATED DOCUMENTATION

[Remote Fault Detection for Link Fault Management](#) | [678](#)

[Remote Loopback for Link Fault Management](#) | [680](#)

# Configuring Link Fault Management

IN THIS SECTION

- [Configuring Link Discovery](#) | [661](#)
- [Configuring the OAM PDU Interval](#) | [662](#)
- [Configuring the OAM PDU Threshold](#) | [662](#)
- [Configuring Threshold Values for Local Fault Events on an Interface](#) | [662](#)
- [Disabling the Sending of Link Event TLVs](#) | [663](#)
- [Example: Configuring IEEE 802.3ah OAM Support on an Interface](#) | [663](#)
- [Example: Configuring IEEE 802.3ah OAM Support for an Interface on ACX Series](#) | [664](#)
- [Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge](#) | [667](#)
- [Example: Configuring Ethernet LFM for CCC](#) | [669](#)
- [Example: Configuring Ethernet LFM for Aggregated Ethernet](#) | [670](#)



- [Configuring an OAM Action Profile | 673](#)
- [Specifying the Actions to Be Taken for Link-Fault Management Events | 674](#)
- [Monitoring the Loss of Link Adjacency | 675](#)
- [Monitoring Protocol Status | 676](#)
- [Configuring Threshold Values for Fault Events in an Action Profile | 677](#)
- [Applying an Action Profile | 677](#)

Use this topic to understand how to configure link fault management features on your device. You can also use this topic to configure an action profile to specify the LFM action that must be performed when a specific LFM event occurs and apply the action profile.

## Configuring Link Discovery

When the IEEE 802.3ah OAM protocol is enabled on a physical interface, the discovery process is automatically triggered. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard.

You can specify the discovery mode used for IEEE 802.3ah OAM support. The discovery process is triggered automatically when OAM IEEE 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

To configure the discovery mode, include the **link-discovery** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
link-discovery (active | passive);
```

In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery.

SEE ALSO

| [link-discovery | 1226](#)

## Configuring the OAM PDU Interval

Periodic OAM PDUs are sent to perform link monitoring.

You can specify the periodic OAM PDU sending interval for fault detection.

To configure the sending interval, include the **pdu-interval** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
pdu-interval interval;
```

The periodic OAM PDU interval range is from 100 through 1000 milliseconds. The default sending interval is 1000 milliseconds.

SEE ALSO

| [pdu-interval](#) | 1307

## Configuring the OAM PDU Threshold

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

To configure the number of PDUs that can be missed from the peer, include the **pdu-threshold** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
pdu-threshold threshold-value;
```

The threshold value range is from 3 through 10. The default is three PDUs.

SEE ALSO

| [pdu-threshold](#) | 1308

## Configuring Threshold Values for Local Fault Events on an Interface

You can configure threshold values on an interface for the local errors that trigger the sending of link event TLVs.

To set the error threshold values for sending event TLVs, include the **frame-error**, **frame-period**, **frame-period-summary**, and **symbol-period** statements at the [edit protocols oam ethernet link-fault-management interface *interface-name* event-thresholds] hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
event-thresholds {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

#### SEE ALSO

[event-thresholds | 1175](#)

[frame-error | 1191](#)

[frame-period | 1192](#)

[frame-period-summary | 1193](#)

[symbol-period | 1372](#)

### Disabling the Sending of Link Event TLVs

You can disable the sending of link event TLVs.

To disable the monitoring and sending of PDUs containing link event TLVs in periodic PDUs, include the **no-allow-link-events** statement at the [edit protocols oam ethernet link-fault-management interface *interface-name* negotiation-options] hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name negotiation-options]
no-allow-link-events;
```

#### SEE ALSO

[no-allow-link-events | 1284](#)

### Example: Configuring IEEE 802.3ah OAM Support on an Interface

Configure 802.3ah OAM support on a 10-Gigabit Ethernet interface:

```
[edit]
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface xe-0/0/0 {
          link-discovery active;
          pdu-interval 800;
          pdu-threshold 4;
          remote-loopback;
          negotiation-options {
            allow-remote-loopback;
          }
          event-thresholds {
            frame-error 30;
            frame-period 50;
            frame-period summary 40;
            symbol-period 20;
          }
        }
      }
    }
  }
}
```

SEE ALSO

| [link-fault-management](#) | [1230](#)

### Example: Configuring IEEE 802.3ah OAM Support for an Interface on ACX Series

Junos OS for ACX Series routers allows the Ethernet interfaces on these routers to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

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

#### Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2 or later for ACX Series routers.

- An ACX1000 or ACX2000 router.

### Overview and Topology

In this example, you configure a 10-Gigabit Ethernet interface on an ACX Series router with 802.3ah OAM support, which includes: link discovery, protocol data units (PDUs), remote loopback, negotiation, and event thresholds.

### Configuring IEEE 802.3ah OAM on an ACX Series Router

#### CLI Quick Configuration

To quickly configure IEEE 802.3ah Ethernet OAM, copy the following commands and paste them into the CLI:

```
edit
edit protocols oam ethernet link-fault-management
set interface xe-0/0/0 link-discovery active pdu-interval 800 pdu-threshold 4 remote-loopback negotiation-options
allow-remote-loopback
set interface xe-0/0/0 event-thresholds frame-error 30 frame-period 50 frame-period-summary 40 symbol-period
20
```

#### Step-by-Step Procedure

To configure IEEE 802.3ah OAM support on an interface:

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

```
[edit protocols oam ethernet link-fault-management]
```

```
user@router1# set interface (OAM Link-Fault Management) xe-0/0/0
```

2. Specify that the interface initiates the discovery process by setting the link discovery mode to **active**:

```
user@router# set interface xe-0/0/0 link-discovery active
```

3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800:

```
user@router# set interface xe-0/0/0 pdu-interval 800
```

4. Define the number of OAM PDUs to miss before an error is logged as 4:

```
user@router# set interface xe-0/0/0 pdu-threshold 4
```

5. Configure the remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes:

```
user@router# set interface xe-0/0/0 remote-loopback
```

6. Configure remote loopback support for the local interface:

```
user@router# set interface xe-0/0/0 negotiation-options allow-remote-loopback
```

7. Set the threshold count for sending frame error events to 30:

```
user@router# set interface xe-0/0/0 event-thresholds frame-error 30
```

8. Set the threshold count for sending frame period error events to 50:

```
user@router# set interface xe-0/0/0 event-thresholds frame-period 50
```

9. Configure the threshold count for sending frame period summary error events to 40:

```
user@router# set interface xe-0/0/0 event-thresholds frame-period-summary 40
```

10. Set the threshold count for sending symbol period events to 20:

```
user@router# set interface xe-0/0/0 event-thresholds symbol-period 20
```

## Results

Check the results of the configuration:

```
[edit]
```

```
user@router# show
```

```
[edit]
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface xe-0/0/0 {
          link-discovery active;
          pdu-interval 800;
          pdu-threshold 4;
          remote-loopback;
          negotiation-options {
            allow-remote-loopback;
          }
          event-thresholds {
            frame-error 30;
            frame-period 50;
            frame-period-summary 40;
            symbol-period 20;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}
}
}
}

```

SEE ALSO


[link-fault-management](#) | 1230

### Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge

In this example, LFM is enabled on an IP link between the provider edge (PE) and customer edge (CE) interfaces. If the link goes down, the fault will be detected by LFM and the interfaces on both sides will be marked **Link-Layer-Down**. This results in notifications to various subsystems (for example, routing) which will take appropriate action.

The link running LFM is shown in [Figure 20 on page 667](#).

Figure 20: Ethernet LFM Between Provider Edge and Customer Edge



To configure Ethernet LFM on an IP link between PE and CE interfaces:

1. Configure LFM on the PE router:

```

[edit]
interfaces ge-1/1/0 {
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
  }
}
protocols {
  oam {

```

```

    ethernet {
        link-fault-management {
            interface ge-1/1/0 {
                pdu-interval 1000;
                pdu-threshold 5;
            }
        }
    }
}

```

## 2. Configure LFM on the CE router:

```

[edit]
interfaces ge-1/1/0 {
    unit 0 {
        family inet {
            address 11.11.11.2/24;
        }
    }
}
protocols {
    oam {
        ethernet {
            link-fault-management {
                interface ge-1/1/0 {
                    pdu-interval 1000;
                    pdu-threshold 5;
                }
            }
        }
    }
}

```

## SEE ALSO

*Ethernet Interfaces User Guide for Routing Devices*

[IEEE 802.3ah OAM Link-Fault Management Overview | 655](#)

[Example: Configuring Ethernet LFM for CCC | 669](#)

[Example: Configuring Ethernet LFM for Aggregated Ethernet | 670](#)




### Example: Configuring Ethernet LFM for CCC

In this example, LFM is configured between two PEs (PE1 and PE2) connected using CCC. With LFM in place, a link fault will be detected immediately, instead of depending on routing protocols to find the fault on end-to-end CCC connection. This also helps in detecting the exact failed link instead of only finding that the end-to-end CCC connectivity has failed. Also, because LFM runs at the link-layer level, it does not need a IP address to operate and so can be used where bidirectional fault detection (BFD) cannot.

The links running LFM are shown in [Figure 21 on page 669](#)

Figure 21: Ethernet LFM for CCC



To configure Ethernet LFM between two PEs connected using CCC:

1. Configure LFM on the PE1 router with CCC:

```
[edit]
interfaces ge-1/1/0 {
  encapsulation ethernet-ccc;
  unit 0;
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/1/0 {
          pdu-interval 1000;
          pdu-threshold 5;
        }
      }
    }
  }
}
```

## 2. Configure LFM on the PE2 router with CCC:

```
[edit]
interfaces ge-1/0/0 {
  encapsulation ethernet-ccc;
  unit 0;
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/0/0 {
          pdu-interval 1000;
          pdu-threshold 5;
        }
      }
    }
  }
}
```

## SEE ALSO

*Ethernet Interfaces User Guide for Routing Devices*

[IEEE 802.3ah OAM Link-Fault Management Overview | 655](#)

[Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge | 667](#)

[Example: Configuring Ethernet LFM for Aggregated Ethernet | 670](#)


[Example: Configuring Ethernet LFM with Loopback Support | 686](#)

## Example: Configuring Ethernet LFM for Aggregated Ethernet

In this example, LFM is configured on an aggregated Ethernet interface (AE0) between Router 1 and Router 2. When configured on aggregated Ethernet, LFM runs on all the individual member links. LFM is enabled or disabled on the member links as they are added or deleted from the aggregation group. The status of individual links is used to determine the status of the aggregated interface.

The use of LFM with aggregated Ethernet is shown in [Figure 22 on page 671](#).

Figure 22: Ethernet LFM for Aggregated Ethernet



To configure LFM on an aggregated Ethernet interface between two routers:

1. Configure LFM on Router 1 for AE0:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
interfaces ge-1/0/1 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ge-2/0/0 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ae0 {
  unit 0 {
    family inet {
      address 11.11.11.2/24;
    }
  }
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ae0;
      }
    }
  }
}
```

```
}
```

## 2. Configure LFM on Router 2 for AE0:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
interfaces ge-1/0/0 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ge-5/0/0 {
  gigether-options {
    802.3ad ae0;
  }
}
interfaces ae0 {
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
  }
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ae0;
      }
    }
  }
}
```

SEE ALSO

*Ethernet Interfaces User Guide for Routing Devices*

[IEEE 802.3ah OAM Link-Fault Management Overview | 655](#)

[Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge | 667](#)

[Example: Configuring Ethernet LFM for CCC | 669](#)

[Example: Configuring Ethernet LFM with Loopback Support | 686](#)

## Configuring an OAM Action Profile

You can create an action profile to define event fault flags and thresholds and the action to be taken. You can then apply the action profile to one or more interfaces.

To configure an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

```
action-profile profile-name {
  action {
    syslog;
    link-down;
    send-critical-event;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    protocol-down;
  }
}
```

**NOTE:** Starting from Junos OS Release 14.2, whenever link-fault management (LFM) with an action profile is configured to mark the interface as down (by including the link-down statement at the [edit protocols oam ethernet link-fault-management] hierarchy level), the port is placed in the blocked state (STP state). In such a state of the interface, data traffic is not transmitted out on that interface. Because the connectivity-fault management (CFM) downstream maintenance MEPs come up on blocked ports, the CFM sessions come up properly. However, the interface is down and the interface status TLV does not contain the correct status. Only if you configure the port status TLV, the actual status of the port is reflected. The interface status TLV does not carry the actual state of the port.

## SEE ALSO

[Setting a Remote Interface into Loopback Mode | 680](#)

[Enabling Remote Loopback Support on the Local Interface | 681](#)

## Specifying the Actions to Be Taken for Link-Fault Management Events

You can specify the action to be taken by the system when the configured link-fault event occurs. Multiple action profiles can be applied to a single interface. For each action-profile, at least one event and one action must be specified. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all the actions are executed.

You might want to set a lower threshold for a specific action such as logging the error and set a higher threshold for another action such as sending a critical event TLV.

To specify the action, include the **action** statement at the [edit protocols oam ethernet link-fault-management action-profile *profile-name*] hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
event {
    link-adjacency-loss;
    protocol-down;
}
action {
    syslog;
    link-down;
    send-critical-event;
}
```

To create a system log entry when the link-fault event occurs, include the **syslog** statement.

To administratively disable the link when the link-fault event occurs, include the **link-down** statement.

To send IEEE 802.3ah link event TLVs in the OAM PDU when a link-fault event occurs, include the **send-critical-event** statement.

**NOTE:** If multiple actions are specified in the action profile, all of the actions are executed in no particular order.

#### SEE ALSO

[action | 1092](#)

[syslog | 1374](#)

[link-down | 1228](#)

[send-critical-event | 1347](#)

### Monitoring the Loss of Link Adjacency

You can specify actions be taken when link adjacency is lost. When link adjacency is lost, the system takes the action defined in the **action** statement of the action profile.

To configure the system to take action when link adjacency is lost, include the **link-adjacency-loss** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]  
  link-adjacency-loss;
```

#### SEE ALSO

[link-adjacency-loss | 1225](#)

[Enabling Remote Loopback Support on the Local Interface | 681](#)

## Monitoring Protocol Status

The CCC-DOWN flag is associated with a circuit cross-connect (CCC) connection, Layer 2 circuit, and Layer 2 VPN, which send the CCC-DOWN status to the kernel. The CCC-DOWN flag indicates that the CCC is down. The CCC-DOWN status is sent to the kernel when the CCC connection, Layer 2 circuit, or Layer 2 VPN is down. This in turn, brings down the CE-facing PE interface associated with the CCC connection, Layer 2 circuit, or Layer 2 VPN.

When the CCC-DOWN flag is signaled to the IEEE 802.3ah protocol, the system takes the action defined in the **action** statement of the action profile. For additional information about Layer 2 circuits, see the Junos OS Layer 2 Circuits User Guide, Junos OS VPNs Configuration Guide.

To monitor the IEEE 802.3ah protocol, on the CE-facing PE interface, include the **protocol-down** statement at the `[edit protocols oam ethernet link-fault-management action-profile profile-name event]` hierarchy level:

1. In configuration mode, go to the `[edit protocols oam ethernet link-fault-management action-profile profile-name event]` hierarchy level.

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

2. Include the **protocol-down** statement.

```
[edit protocols oam ethernet link-fault-management action-profile profile-name event]
user@host# set protocol-down
```

**NOTE:** If multiple events are specified in the action profile, all the events must occur before the specified action is taken.

## SEE ALSO

[protocol-down](#) | 1328

[Setting a Remote Interface into Loopback Mode](#) | 680

[Enabling Remote Loopback Support on the Local Interface](#) | 681



### Configuring Threshold Values for Fault Events in an Action Profile

You can configure link event thresholds for received error events that trigger the action specified in the **action** statement. You can then apply the action profile to one or more interfaces.

To configure link event thresholds, include the **link-event-rate** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

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

SEE ALSO

| [link-event-rate](#) | [1229](#)

### Applying an Action Profile

You can apply an action profile to one or more interfaces.

To apply an action profile to an interface, include the **apply-action-profile** statement at the **[edit protocols oam ethernet link-fault-management action-profile interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
  apply-action-profile profile-name;
```

SEE ALSO

| [apply-action-profile](#) | [1112](#)

### Release History Table

| Release              | Description                         |
|----------------------|-------------------------------------|
| <a href="#">14.2</a> | Starting from Junos OS Release 14.2 |

## RELATED DOCUMENTATION

[Remote Loopback for Link Fault Management | 680](#)

## Remote Fault Detection for Link Fault Management

### IN THIS SECTION

- [Detecting Remote Faults | 678](#)
- [Enabling Dying Gasp Functionality | 679](#)

Use this topic to understand more about remote faults and how they are detected and also how to enable the dying gasp feature to avoid file system corruption for LFM.

### Detecting Remote Faults

Fault detection is either based on flags or fault event type, length, and values (TLVs) received in OAM protocol data units (PDUs). Flags that trigger a link fault are:

- Critical Event
- Dying Gasp
- Link Fault

The link event TLVs are sent by the remote DTE by means of event notification PDUs. Link event TLVs are:

- Errored Symbol Period Event
- Errored Frame Event
- Errored Frame Period Event
- Errored Frame Seconds Summary Event

### SEE ALSO

[IEEE 802.3ah OAM Link-Fault Management Overview | 655](#)[Configuring IEEE 802.3ah OAM Link-Fault Management](#)

## Enabling Dying Gasp Functionality

Dying gasp means an unrecoverable condition such as a power failure. In this condition, the local peer informs the remote peer about the failure state. When the remote peer receives a dying-gasp PDU, it takes an action corresponding to the action profile configured with the **link-adjacency-loss** event. Dying gasp helps to avoid file system corruption.

**NOTE:** ACX5096 and ACX5048 routers do not support dying-gasp.

ACX Series routers can generate and receive dying-gasp packets. When LFM is configured on an interface, a dying-gasp PDU is generated for the interface on the following failure conditions:

- Power failure
- Packet Forwarding Engine panic or a crash

ACX Series routers support the following CLI statements to enable dying-gasp functionality:

- **dgasp-int**—Enables dying-gasp functionality.
- **dgasp-usb**—Resets USB port during dying-gasp event.

The **dgasp-int** and **dgasp-usb** CLI statements are added under the **[edit system]** hierarchy to enable dying-gasp functionality.

To enable dying-gasp functionality, you need to configure the **dgasp-int** and **dgasp-usb** CLI statements as shown below:

```
root@host% cli
root@host> configure
Entering configuration mode

[edit]
root@host# set system dgasp-int

[edit]
root@host# set system dgasp-usb

[edit]
root@host# commit

commit complete

[edit]
root@host# show system
```

```
dgasp-int;
dgasp-usb;
```

The dying-gasp functionality is disabled by default.

#### SEE ALSO

[Understanding Ethernet OAM Link Fault Management for ACX Series Routers](#) | 656

#### RELATED DOCUMENTATION

[Introduction to OAM Link Fault Management \(LFM\)](#) | 655

## Remote Loopback for Link Fault Management

### IN THIS SECTION

- [Setting a Remote Interface into Loopback Mode](#) | 680
- [Enabling Remote Loopback Support on the Local Interface](#) | 681
- [Enabling Nonstop Routing for Ethernet Link Fault Management on Backup Routers](#) | 682
- [Example: Configuring Ethernet LFM with Loopback Support](#) | 686

Use this topic to understand what happens when you set a remote interfaces in loopback mode and how to enable remote loopback. You can also learn how to enable nonstop routing for LFM.

### Setting a Remote Interface into Loopback Mode

You can configure the software to set the remote DTE into loopback mode on the following interfaces:

- IQ2 and IQ2-E Gigabit Ethernet interfaces
- Ethernet interfaces on the MX Series routers or EX Series switches

Junos OS can place a remote DTE into loopback mode (if remote-loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote-loopback

request and puts the interface into remote-loopback mode. When the interface is in remote-loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent to the management plane and processed.

To configure remote loopback, include the **remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
remote-loopback;
```

To take the remote DTE out of loopback mode, remove the **remote-loopback** statement from the configuration.

SEE ALSO

| [remote-loopback](#) | 1337

## Enabling Remote Loopback Support on the Local Interface

You can allow a remote DTE to set a local interface into remote loopback mode on IQ2 and IQ2-E Gigabit Ethernet interfaces and all Ethernet interfaces on the MX Series routers and EX Series switches. When a remote-loopback request is sent by a remote DTE, the Junos OS places the local interface into loopback mode. When an interface is in loopback mode, all frames except OAM PDUs are looped back without any changes to the frames. OAM PDUs continue to be sent to the management plane and processed. By default, the remote loopback feature is not enabled.

To enable remote loopback, include the **allow-remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name* negotiation-options]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name negotiation-options]  
allow-remote-loopback;
```

**NOTE:** Activation of OAM remote loopback may result in data frame loss.

SEE ALSO

| [allow-remote-loopback](#) | 1111

## Enabling Nonstop Routing for Ethernet Link Fault Management on Backup Routers

Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured. When the lfmd daemon runs on the backup Routing Engine as well, the link fault management states are kept in sync and so minimal effort is required by the lfmd daemon post switch over.

To enable Nonstop routing for Ethernet LFM on backup routers:

1. Enable graceful Routing Engine switchover. By default, GRES is disabled. To enable GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. By default, Nonstop routing is disabled. When you enable GRES, NSR is enabled.

```
[edit chassis redundancy]
user@host# set graceful-switchover
```

2. Synchronize the Routing Engine configuration. To synchronize the master Routing Engine configuration with the backup, include the **synchronize** statement at the **[edit system]** hierarchy level.

```
[edit system]
user@host# set commit synchronize
```

3. After enabling nonstop routing, commit the configuration.

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

4. To verify if nonstop routing is enabled on the backup router, at the operational mode, use the **show oam ethernet link-fault-management** command on the master router and then the backup router. Because you have enabled synchronization, the output of the master router and the backup router is identical. However, the statistics maintained by the master router are not synchronized with the backup router..

```
{master}
user@host# show oam ethernet link-fault-management ge-0/2/0 detail
```

```
Interface: ge-0/2/0
  Status: Running, Discovery state: Send Any
  Transmit interval: 100ms, PDU threshold: 3 frames, Hold time: 300ms
  Peer address: ac:4b:c8:81:90:a4
```

```
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
    Information: 0, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 0, Organization specific: 0
OAM flags receive statistics:
    Critical event: 0, Dying gasp: 0, Link fault: 0
OAM transmit statistics:
    Information: 0, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 786, Organization specific: 0
OAM received symbol error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
OAM received frame error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
OAM received frame period error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
OAM received frame seconds error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
OAM current symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
OAM current frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
Loopback tracking: Enabled, Loop status: Not Found
Detect LOC: Enabled, LOC status: Not Found
Remote entity information:
    Remote MUX action: forwarding, Remote parser action: forwarding
    Discovery mode: active, Unidirectional mode: unsupported
    Remote loopback mode: unsupported, Link events: supported
    Variable requests: unsupported
Application profile statistics:
    Profile Name                Invoked      Executed
    LK_ADJ_LOSS100_1            1            1
    LK_ADJ_LOSS100_2            1            0
```

|                  |   |   |
|------------------|---|---|
| LK_ADJ_LOSS100_3 | 1 | 0 |
| LK_ADJ_LOSS101_1 | 1 | 1 |
| LK_ADJ_LOSS101_2 | 1 | 0 |
| LK_ADJ_LOSS101_3 | 1 | 0 |
| LK_ADJ_LOSS106_1 | 0 | 0 |
| LK_ADJ_LOSS106_2 | 0 | 0 |
| LK_ADJ_LOSS106_3 | 0 | 0 |
| LK_ADJ_LOSS107_1 | 0 | 0 |
| LK_ADJ_LOSS107_2 | 0 | 0 |
| LK_ADJ_LOSS107_3 | 0 | 0 |

{backup}

user@host# show oam ethernet link-fault-management ge-0/2/0 detail

Interface: ge-0/2/0

Status: Running, Discovery state: Send Any

Transmit interval: 100ms, PDU threshold: 3 frames, Hold time: 300ms

Peer address: ac:4b:c8:81:90:a4

Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50

OAM receive statistics:

Information: 0, Event: 0, Variable request: 0, Variable response: 0

Loopback control: 0, Organization specific: 0

OAM flags receive statistics:

Critical event: 0, Dying gasp: 0, Link fault: 0

OAM transmit statistics:

Information: 0, Event: 0, Variable request: 0, Variable response: 0

Loopback control: 786, Organization specific: 0

OAM received symbol error event information:

Events: 0, Window: 0, Threshold: 0

Errors in period: 0, Total errors: 0

OAM received frame error event information:

Events: 0, Window: 0, Threshold: 0

Errors in period: 0, Total errors: 0

OAM received frame period error event information:

Events: 0, Window: 0, Threshold: 0

Errors in period: 0, Total errors: 0

OAM received frame seconds error event information:

Events: 0, Window: 0, Threshold: 0

Errors in period: 0, Total errors: 0

OAM transmitted symbol error event information:

Events: 0, Window: 0, Threshold: 1

Errors in period: 0, Total errors: 0



```
OAM current symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM current frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
Loopback tracking: Enabled, Loop status: Not Found
Detect LOC: Enabled, LOC status: Not Found
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: unsupported, Link events: supported
  Variable requests: unsupported
Application profile statistics:
  Profile Name                Invoked      Executed
LK_ADJ_LOSS100_1              0            0
LK_ADJ_LOSS100_2              0            0
LK_ADJ_LOSS100_3              0            0
LK_ADJ_LOSS101_1              0            0
LK_ADJ_LOSS101_2              0            0
LK_ADJ_LOSS101_3              0            0
LK_ADJ_LOSS106_1              0            0
LK_ADJ_LOSS106_2              0            0
LK_ADJ_LOSS106_3              0            0
LK_ADJ_LOSS107_1              0            0
LK_ADJ_LOSS107_2              0            0
LK_ADJ_LOSS107_3              0            0
```

**NOTE:** After the switchover, if issues are observed, use the **clear oam ethernet link-fault-management state** command for specific sessions. If the issue does not get resolved, restart the lfmd daemon.


SEE ALSO

|                                                                       |
|-----------------------------------------------------------------------|
| <a href="#">IEEE 802.3ah OAM Link-Fault Management Overview   655</a> |
| <a href="#">show oam ethernet link-fault-management   1962</a>        |

## Example: Configuring Ethernet LFM with Loopback Support

In this example, LFM is configured between provider edge (PE) router and the customer edge (CE) router. The PE router can put the CE router in remote loopback mode. This allows the PE to have all the traffic sent to the CE router looped back for diagnostics purposes, as shown in [Figure 23 on page 686](#).

Figure 23: Ethernet LFM with Loopback Support



To configure LFM between a PE router and a CE router:

1. Configure LFM loopback on the PE router:

```
[edit]
interfaces ge-1/0/0 {
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
  }
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/0/0 {
          pdu-interval 1000;
          pdu-threshold 5;
          remote-loopback;
        }
      }
    }
  }
}
```

2. Configure LFM loopback on the CE router:

```
[edit]
```

```

interfaces ge-1/1/0 {
  unit 0 {
    family inet {
      address 11.11.11.2/24;
    }
  }
}
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-1/1/0 {
          pdu-interval 1000;
          pdu-threshold 5;
          negotiation-options {
            allow-remote-loopback;
          }
        }
      }
    }
  }
}

```

**NOTE:** If the **negotiation options allow-remote-loopback** statement on the CE router is deleted before removing the CE router from remote loopback mode, traffic flow between the PE router and CE router is affected. Hence, delete the **remote-loopback** statement on the PE router before deleting the **negotiation-options allow-remote-loopback** statement on the CE router.

## SEE ALSO

[Example: Configuring Ethernet LFM Between Provider Edge and Customer Edge | 667](#)

[Example: Configuring Ethernet LFM for CCC | 669](#)

[Example: Configuring Ethernet LFM for Aggregated Ethernet | 670](#)

Release History Table

| Release                | Description                                                                                                                                                                                   |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">17.3R1</a> | Starting in Junos OS Release 17.3R1, the Ethernet link fault management daemon (lfmd) runs on the backup Routing Engine as well when graceful Routing Engine switchover (GRES) is configured. |

RELATED DOCUMENTATION

|                                                                       |
|-----------------------------------------------------------------------|
| <a href="#">Introduction to OAM Link Fault Management (LFM)   655</a> |
| <a href="#">Configuring Link Fault Management   660</a>               |

# Configuring ITU-TY.1731 Ethernet Service OAM

## IN THIS CHAPTER

- [ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)
- [Configuring Ethernet Frame Delay Measurement Sessions | 709](#)
- [Configuring Ethernet Frame Loss Measurement | 747](#)
- [Configuring an Iterator Profile | 781](#)
- [Configuring Ethernet Synthetic Loss Measurements | 796](#)
- [Ethernet Alarm Indication | 810](#)
- [Inline Transmission Mode | 824](#)

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

### IN THIS SECTION

- [Ethernet Frame Delay Measurements Overview | 690](#)
- [Ethernet Frame Loss Measurement Overview | 696](#)
- [Service-Level Agreement Measurement | 698](#)
- [On-Demand Mode for SLA Measurement | 699](#)
- [Proactive Mode for SLA Measurement | 700](#)
- [Ethernet Failure Notification Protocol Overview | 701](#)
- [Ethernet Synthetic Loss Measurement Overview | 702](#)
- [Scenarios for Configuration of ETH-SLM | 703](#)
- [Format of ETH-SLM Messages | 705](#)
- [Transmission of ETH-SLM Messages | 706](#)

Use this topic to understand more about service OAM (ITU-TY.1731) and its two main components: fault management (monitoring, detection, and isolation) and performance monitoring (frame loss measurement, synthetic frame loss measurement, and frame delay measurement).

## Ethernet Frame Delay Measurements Overview

### IN THIS SECTION

- [ITU-T Y.1731 Frame Delay Measurement Feature | 690](#)
- [One-Way Ethernet Frame Delay Measurement | 692](#)
- [Two-Way Ethernet Frame Delay Measurement | 693](#)
- [Choosing Between One-Way and Two-Way ETH-DM | 695](#)
- [Restrictions for Ethernet Frame Delay Measurement | 695](#)

### ***ITU-T Y.1731 Frame Delay Measurement Feature***

The IEEE 802.3-2005 standard for Ethernet Operations, Administration, and Maintenance (OAM) defines a set of link fault management mechanisms to detect and report link faults on a single point-to-point Ethernet LAN.

Junos OS supports key OAM standards that provide for automated end-to-end management and monitoring of Ethernet service by service providers:

- *IEEE Standard 802.1ag*, also known as “Connectivity Fault Management (CFM).”
- *ITU-T Recommendation Y.1731*, which uses different terminology than IEEE 802.1ag and defines Ethernet service OAM features for fault monitoring, diagnostics, and performance monitoring.

These capabilities allow operators to offer binding service-level agreements (SLAs) and generate new revenues from rate- and performance-guaranteed service packages that are tailored to the specific needs of their customers.

ACX Series routers support proactive and on-demand modes.

**NOTE:** ACX5048 and ACX5096 routers supports only software-based time stamping for delay measurement.

## Ethernet CFM

The IEEE 802.1ag standard for connectivity fault management (CFM) defines mechanisms to provide for end-to-end Ethernet service assurance over any path, whether a single link or multiple links spanning networks composed of multiple LANs.

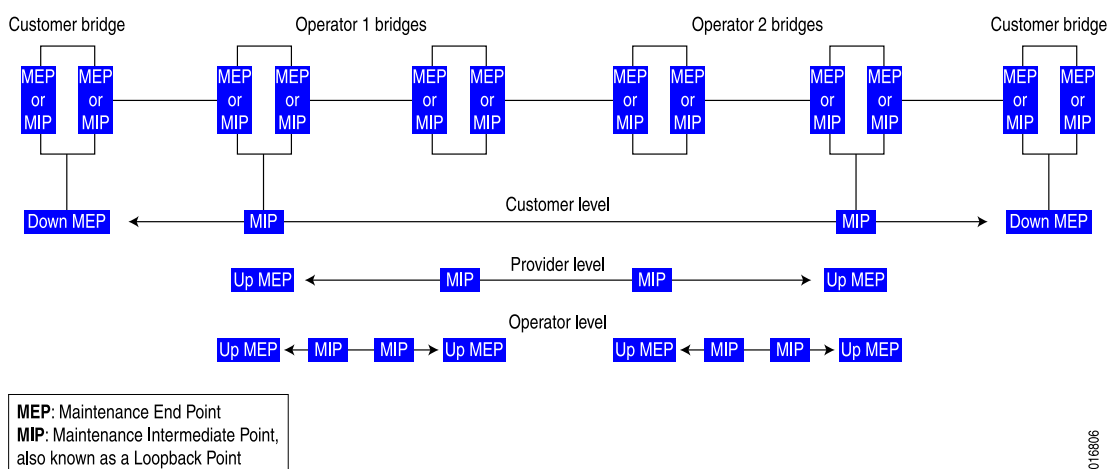
For Ethernet interfaces on M320, MX Series, and T Series routers, Junos OS supports the following key elements of the Ethernet CFM standard:

- Fault monitoring using the IEEE 802.1ag Ethernet OAM Continuity Check protocol
- Path discovery and fault verification using the IEEE 802.1ag Ethernet OAM Linktrace protocol
- Fault isolation using the IEEE 802.1ag Ethernet OAM Loopback protocol

In a CFM environment, network entities such as network operators, service providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management and end-to-end monitoring, and still avoid security breaches.

Figure 24 on page 691 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

**Figure 24: Relationship of MEPs, MIPs, and Maintenance Domain Levels**



**NOTE:** On ACX Series routers, the maintenance intermediate points (MIP) is supported only on the ACX5048 and ACX5096 routers.

### **Ethernet Frame Delay Measurement**

Two key objectives of OAM functionality are to measure quality-of-service attributes such as frame delay and frame delay variation (also known as “frame jitter”). Such measurements can enable you to identify network problems before customers are impacted by network defects.

Junos OS supports Ethernet frame delay measurement between MEPs configured on Ethernet physical or logical interfaces on MX Series routers. Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor SLAs. Ethernet frame delay measurement also collects other useful information, such as worst and best case delays, average delay, and average delay variation. The Junos OS implementation of Ethernet frame delay measurement (ETH-DM) is fully compliant with the ITU-T Recommendation Y.1731, *OAM Functions and Mechanisms for Ethernet-based Networks*. The recommendation defines OAM mechanisms for operating and maintaining the network at the Ethernet service layer, which is called the “ETH layer” in ITU-T terminology.

MX Series routers with modular port concentrators (MPCs) and 10-Gigabit Ethernet MPCs with SFP+ support ITU-T Y.1731 functionality on VPLS for frame-delay and delay-variation.

**NOTE:** MX Series Virtual Chassis does not support Ethernet frame delay measurement (DM).

### **One-Way Ethernet Frame Delay Measurement**

In one-way ETH-DM mode, a series of frame delay and frame delay variation values are calculated based on the time elapsed between the time a measurement frame is sent from the initiator MEP at one router and the time when the frame is received at the receiver MEP at the other router.

**NOTE:** ACX Series routers do not support one-way Ethernet frame delay measurement.

#### **1DM Transmission**

When you start a one-way frame delay measurement, the router sends 1DM frames—frames that carry the protocol data unit (PDU) for a one-way delay measurement—from the initiator MEP to the receiver MEP at the rate and for the number of frames you specify. The router marks each 1DM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

#### **1DM Reception**

When an MEP receives a 1DM frame, the router that contains the receiver MEP measures the one-way delay for that frame (the difference between the time the frame was received and the timestamp contained in the frame itself) and the delay variation (the difference between the current and previous delay values).



### ***One-Way ETH-DM Statistics***

The router that contains the receiver MEP stores each set of one-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

### ***One-Way ETH-DM Frame Counts***

Each router counts the number of one-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number of 1DM frames sent.
- For a receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.

### ***Synchronization of System Clocks***

The accuracy of one-way delay calculations depends on close synchronization of the system clocks at the initiator MEP and receiver MEP.

The accuracy of one-way delay variation is not dependent on system clock synchronization. Because delay variation is simply the difference between consecutive one-way delay values, the out-of-phase period is eliminated from the frame jitter values.

**NOTE:** For a given one-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only on the router that contains the receiver MEP.

### ***Two-Way Ethernet Frame Delay Measurement***

In two-way ETH-DM mode, frame delay and frame delay variation values are based on the time difference between when the initiator MEP transmits a request frame and receives a reply frame from the responder MEP, subtracting the time elapsed at the responder MEP.

### ***DMM Transmission***

When you start a two-way frame delay measurement, the router sends delay measurement message (DMM) frames— frames that carry the PDU for a two-way ETH-DM request—from the initiator MEP to the responder MEP at the rate and for the number of frames you specify. The router marks each DMM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

### DMR Transmission

When an MEP receives a DMM frame, the responder MEP responds with a delay measurement reply (DMR) frame, which carries ETH-DM reply information and a copy of the timestamp contained in the DMM frame.

### DMR Reception

When an MEP receives a valid DMR, the router that contains the MEP measures the two-way delay for that frame based on the following sequence of timestamps:

1.  $TI_{TxDMM}$
2.  $TR_{RxDMM}$
3.  $TR_{TxDMR}$
4.  $TI_{RxDMR}$

A two-way frame delay is calculated as follows:

$$[TI_{RxDMR} - TI_{TxDMM}] - [TR_{TxDMR} - TR_{RxDMM}]$$

The calculation shows that frame delay is the difference between the time at which the initiator MEP sends a DMM frame and the time at which the initiator MEP receives the associated DMR frame from the responder MEP, minus the time elapsed at the responder MEP.

The delay variation is the difference between the current and previous delay values.

### Two-Way ETH-DM Statistics

The router that contains the initiator MEP stores each set of two-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

### Two-Way ETH-DM Frame Counts

Each router counts the number of two-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number DMM frames transmitted, the number of valid DMR frames received, and the number of invalid DMR frames received.
- For a responder MEP, the router counts the number of DMR frames sent.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.

**NOTE:** For a given two-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only at the router that contains the initiator MEP.

### **Choosing Between One-Way and Two-Way ETH-DM**

One-way frame delay measurement requires that the system clocks at the initiator MEP and receiver MEP are closely synchronized. Two-way frame delay measurement does not require synchronization of the two systems. If it is not practical for the clocks to be synchronized, two-way frame delay measurements are more accurate.

When two systems are physically close to each other, their one-way delay values are very high compared to their two-way delay values. One-way delay measurement requires that the timing for the two systems be synchronized at a very granular level, and MX Series routers currently do not support this granular synchronization.

### **Restrictions for Ethernet Frame Delay Measurement**

The following restrictions apply to the Ethernet frame delay measurement feature:

- The ETH-DM feature is not supported on label-switched interface (LSI) pseudowires.  
The ETH-DM feature is supported on aggregated Ethernet interfaces.
- Hardware-assisted timestamping for ETH-DM frames in the reception path is only supported for MEP interfaces on Enhanced DPCs and Enhanced Queuing DPCs in MX Series routers. For information about hardware-assisted timestamping, see [“Guidelines for Configuring Routers to Support an ETH-DM Session” on page 710](#) and [“Enabling the Hardware-Assisted Timestamping Option” on page 722](#).
- Ethernet frame delay measurements can be triggered only when the distributed periodic packet management daemon (**ppm**) is enabled. For more information about this limitation, see [“Guidelines for Configuring Routers to Support an ETH-DM Session” on page 710](#) and [“Ensuring That Distributed ppm Is Not Disabled” on page 719](#).
- You can monitor only one session at a time to the same remote MEP or MAC address. For more information about starting an ETH-DM session, see [“Starting an ETH-DM Session” on page 725](#).
- ETH-DM statistics are collected at only one of the two peer routers in the ETH-DM session. For a one-way ETH-DM session, you can display frame ETH-DM statistics at the receiver MEP only, using ETH-DM-specific **show** commands. For a two-way ETH-DM session, you can display frame delay statistics at the initiator MEP only, using the same ETH-DM-specific **show** commands. For more information, see [“Managing ETH-DM Statistics and ETH-DM Frame Counts” on page 743](#).
- ETH-DM frame counts are collected at both MEPs and are stored in the respective CFM databases.
- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. Therefore, the collection of ETH-DM statistics and ETH-DM frame counters has to be restarted, after the switchover is complete. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *High Availability User Guide*.
- Accuracy of frame delay statistics is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

SEE ALSO

[Ethernet Alarm Indication](#) | 810

[Inline Transmission Mode](#) | 824

## Ethernet Frame Loss Measurement Overview

The key objectives of the OAM functionality are to measure quality-of-service attributes such as frame delay, frame delay variation (also known as “frame jitter”), and frame loss. Such measurements enable you to identify network problems before customers are impacted by network defects. For more information about Ethernet frame delay measurement, see [“Ethernet Frame Delay Measurements Overview” on page 690](#).

Junos OS supports Ethernet frame loss measurement (ETH-LM) between maintenance association end points (MEPs) configured on Ethernet physical or logical interfaces on MX Series routers and is presently supported only for VPWS service. ETH-LM is used by operators to collect counter values applicable for ingress and egress service frames. These counters maintain a count of transmitted and received data frames between a pair of MEPs. Ethernet frame loss measurement is performed by sending frames with ETH-LM information to a peer MEP and similarly receiving frames with ETH-LM information from the peer MEP. This type of frame loss measurement is also known as single-ended Ethernet loss measurement.

**NOTE:** MX Series Virtual Chassis does not support Ethernet frame loss measurement (ETH-LM).

ETH-LM supports the following frame loss measurements:

- Near-end frame loss measurement—Measurement of frame loss associated with ingress data frames.
- Far-end frame loss measurement—Measurement of frame loss associated with egress data frames.

**NOTE:** The proactive and dual-ended loss measurement functionality of ITU-T Y1731 is not supported on the ACX Series routers.

The ETH-LM feature is supported on aggregated Ethernet interfaces.

**NOTE:** Starting Junos OS Release 16.1, the Ethernet loss measurement (ETH-LM) results are inaccurate when connectivity fault management (CFM) and performance monitoring (PM) PDUs received locally at a maintenance endpoint (MEP) as classified as belonging to the yellow class or a packet loss priority (PLP) of medium-high. This problem of incorrect results is specific to Ethernet loss measurement for CFM sessions of down MEPs. The Ethernet loss measurement statistics are inaccurate in the following scenarios:

- Ethernet loss measurement is working on a CFM session for a MEP in down state
- CFM PDUs received on the logical interface of the down MEP are classified by the classifier as yellow or medium-high PLP
- A packet is identified as yellow when the input classifier marks the PLP as medium-high.

The problem of discrepancies with Ethernet loss measurement results is not observed when you configure Ethernet loss measurement in colorless mode. To avoid this problem of inaccurate loss measurement results, provision all local CFM PDUs as green or with the PLP as high.

**NOTE:** 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.

## SEE ALSO

[Managing Continuity Measurement Statistics | 740](#)

[Example: Measuring Ethernet Frame Loss for Single-Tagged LMM/LMR PDUs | 750](#)

[Example: Measuring Ethernet Frame Loss for Dual-Tagged LMM/LMR PDUs | 766](#)

## Service-Level Agreement Measurement

Service-level agreement (SLA) measurement is the process of monitoring the bandwidth, delay, delay variation (jitter), continuity, and availability of a service (E-Line or E-LAN). It enables you to identify network problems before customers are impacted by network defects.

### NOTE:

The Ethernet VPN services can be classified into:

- Peer-to-peer-services (E-Line services)—The E-Line services are offered using MPLS-based Layer 2 VPN virtual private wire service (VPWS).
- Multipoint-to-multipoint services (E-LAN services)—The E-LAN services are offered using MPLS-based virtual private LAN service (VPLS).

For more information, see the *Junos VPNs Configuration Guide*.

In Junos OS, SLA measurements are classified into:

- On-demand mode—In on-demand mode, the measurements are triggered through the CLI. For more information, see [“On-Demand Mode for SLA Measurement” on page 699](#).
- Proactive mode—In proactive mode, the measurements are triggered by an iterator application. For more information, see [“Proactive Mode for SLA Measurement” on page 700](#).

For more information about frame delay measurement, see [“Ethernet Frame Delay Measurements Overview” on page 690](#). For more information about frame loss measurement, see [“Ethernet Frame Loss Measurement Overview” on page 696](#). Note that Ethernet frame delay measurement and Ethernet frame loss measurement are not supported on the **ae** interface.

### SEE ALSO

[Ethernet Alarm Indication | 810](#)

[Inline Transmission Mode | 824](#)

## On-Demand Mode for SLA Measurement

In on-demand mode, the measurements are triggered by the user through the CLI.

When the user triggers the delay measurement through the CLI, the delay measurement request that is generated is as per the frame formats specified by the ITU-T Y.1731 standard. For two-way delay measurement, the server-side processing can be delegated to the Packet Forwarding Engine to prevent overloading on the Routing Engine. For more information, see [“Configuring Routers to Support an ETH-DM Session” on page 718](#). When the server-side processing is delegated to the Packet Forwarding Engine, the delay measurement message (DMM) frame **receive** counters and delay measurement reply (DMR) frame **transmit** counters are not displayed by the **show** command.

When the user triggers the loss measurement through the CLI, the router sends the packets in standard format along with the loss measurement TLV. By default, the **session-id-tlv** argument is included in the packet to allow concurrent loss measurement sessions from same local MEP. You can also disable the session ID TLV by using the **no-session-id-tlv** argument.

Single-ended ETH-LM is used for on-demand operation, administration, and maintenance purposes. An MEP sends frames with ETH-LM request information to its peer MEP and receives frames with ETH-LM reply information from its peer MEP to carry out loss measurements. The protocol data unit (PDU) used for a single-ended ETH-LM request is referred to as a loss measurement message (LMM) and the PDU used for a single-ended ETH-LM reply is referred to as a loss measurement reply (LMR).

### SEE ALSO

---

[Configuring Ethernet Frame Loss Measurement | 747](#)

---

[Configuring an Iterator Profile | 781](#)

---

[Configuring Ethernet Synthetic Loss Measurements | 796](#)

## Proactive Mode for SLA Measurement

In proactive mode, SLA measurements are triggered by an iterator application. An iterator is designed to periodically transmit SLA measurement packets in form of ITU-Y.1731-compliant frames for two-way delay measurement or loss measurement on MX Series routers. This mode differs from on-demand SLA measurement, which is user initiated. The iterator sends periodic delay or loss measurement request packets for each of the connections registered to it. Iterators make sure that measurement cycles do not occur at the same time for the same connection to avoid CPU overload. Junos OS supports proactive mode for VPWS. For an iterator to form a remote adjacency and to become functionally operational, the continuity check message (CCM) must be active between the local and remote MEP configurations of the connectivity fault management (CFM). Any change in the iterator adjacency parameters resets the existing iterator statistics and restarts the iterator. Here, the term adjacency refers to a pairing of two endpoints (either connected directly or virtually) with relevant information for mutual understanding, which is used for subsequent processing. For example, the iterator adjacency refers to the iterator association between the two endpoints of the MEPs.

For every DPC or MPC, only 30 iterator instances for a cycle time value of 10 milliseconds (ms) are supported. In Junos OS, 255 iterator profile configurations and 2000 remote MEP associations are supported.

Iterators with cycle time value less than 100 ms are supported only for infinite iterators, whereas the iterators with cycle time value greater than 100 ms are supported for both finite and infinite iterators. Infinite iterators are iterators that run infinitely until the iterator is disabled or deactivated manually.

**NOTE:** ACX5048 and ACX5096 routers supports iterator cycle time of only 1 second and above.

A VPWS service configured on a router is monitored for SLA measurements by registering the connection (here, the connection is a pair of remote and local MEPs) on an iterator and then initiating periodic SLA measurement frame transmission on those connections. The end-to-end service is identified through a maintenance association end point (MEP) configured at both ends.

For two-way delay measurement and loss measurement, an iterator sends a request message for the connection in the list (if any) and then sends a request message for the connection that was polled in the former iteration cycle. The back-to-back request messages for the SLA measurement frames and their responses help in computing delay variation and loss measurement.

The Y.1731 frame transmission for a service attached to an iterator continues endlessly unless intervened and stopped by an operator or until the iteration-count condition is met. To stop the iterator from sending out any more proactive SLA measurement frames, the operator must perform one of the following tasks:

- Enable the **deactivate sla-iterator-profile** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *md-name* maintenance association *ma-name* mep**



**mep-id remote-mep mep-id]** hierarchy level. For more information, see [“Verifying the Configuration of an Iterator Profile” on page 785](#).

- Provision a **disable** statement under the corresponding iterator profile at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles profile-name]** hierarchy level. For more information, see [“Configuring an Iterator Profile” on page 782](#).

### ***Ethernet Delay Measurements and Loss Measurement by Proactive Mode***

In two-way delay measurement, the delay measurement message (DMM) frame is triggered through an iterator application. The DMM frame carries an iterator type, length, and value (TLV) in addition to the fields described in standard frame format and the server copies the iterator TLV from the DMM frame to the delay measurement reply (DMR) frame.

In one-way delay variation computation using the two-way delay measurement method, the delay variation computation is based on the timestamps that are present in the DMR frame (and not the 1DM frame). Therefore, there is no need for client-side and server-side clocks to be in sync. Assuming that the difference in their clocks remains constant, the one-way delay variation results are expected to be fairly accurate. This method also eliminates the need to send separate 1DM frames just for the one-way delay variation measurement purpose.

In proactive mode for loss measurement, the router sends packets in standard format along with loss measurement TLV and iterator TLV.

### SEE ALSO

[Configuring an Iterator Profile | 782](#)

[Configuring a Remote MEP with an Iterator Profile | 794](#)

[Verifying the Configuration of an Iterator Profile | 785](#)

[Managing Iterator Statistics | 788](#)

## **Ethernet Failure Notification Protocol Overview**

The Failure Notification Protocol (FNP) is a failure notification mechanism that detects failures in Point-to-Point Ethernet transport networks on MX Series routers. If a node link fails, FNP detects the failure and sends out FNP messages to the adjacent nodes that a circuit is down. Upon receiving the FNP message, nodes can redirect traffic to the protection circuit.

**NOTE:** FNP is supported on E-Line services only.

An E-Line service provides a secure Point-to-Point Ethernet connectivity between two user network interfaces (UNIs). E-Line services are a protected service and each service has a working circuit and protection circuit. CFM is used to monitor the working and protect paths. CCM intervals result in failover time in hundreds of milliseconds or a few seconds. FNP provides service circuit failure detection and propagation in less than 50ms and provide 50ms failover for E-Line services.

The MX router acts as a PE node and handles the FNP messages received on the management VLAN and the FNP messages received on both the Ethernet interfaces and PWs created for the management VPLS. MX-series routers do not initiate FNP messages and responds only to FNP messages generated by devices in the Ethernet Access network. FNP can be enabled only on logical interfaces that are part of a VPLS routing instance, and no physical interfaces in that VPLS routing instance should have CCM configured. FNP can be enabled only on one logical interface per physical interface.

All E-Line services are configured as layer 2 circuits with edge protection. A VLAN associated with the working circuit or protection circuit must map to a logical interface. No trunk port or access port is supported in the ring link for VLANs used by E-LINE services. FNP does not control the logical interface associated with protection circuit. Only E-Line service whose termination point is not in an MX node is controlled by FNP.

FNP supports graceful restart and the Graceful Routing Engine switchover (GRES) features.

SEE ALSO

*Configuring the Failure Notification Protocol*

[show oam ethernet fnp interface | 1956](#)

[show oam ethernet fnp status | 1960](#)

[show oam ethernet fnp messages | 1958](#)

[connectivity-fault-management | 1132](#)

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)

## Ethernet Synthetic Loss Measurement Overview

Ethernet synthetic loss measurement (ETH-SLM) is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic. Each maintenance association end point (MEP) performs frame loss measurements, which contribute to unavailable time.

A near-end frame loss specifies frame loss associated with ingress data frames and a far-end frame loss specifies frame loss associated with egress data frames. Both near-end and far-end frame loss measurements contribute to near-end severely errored seconds and far-end severely errored seconds that are used in combination to determine unavailable time. ETH-SLM is performed using synthetic loss message (SLM) and synthetic loss reply (SLR) frames. ETH-SLM facilitates each MEP to perform near-end and far-end

synthetic frame loss measurements by using synthetic frames because a bidirectional service is defined as unavailable if either of the two directions is determined to be unavailable.

There are the two types of frame loss measurement, defined by the ITU-T Y.1731 standards, ETH-LM and ETH-SLM. Junos OS supports only single-ended ETH-SLM. In single-ended ETH-SLM, each MEP sends frames with the ETH-SLM request information to its peer MEP and receives frames with ETH-SLM reply information from its peer MEP to perform synthetic loss measurements. Single-ended ETH-SLM is used for proactive or on-demand OAM to perform synthetic loss measurements applicable to point-to-point Ethernet connection. This method allows a MEP to initiate and report far-end and near-end loss measurements associated with a pair of MEPs that are part of the same maintenance entity group (MEG).

**NOTE:** MX Series Virtual Chassis does not support Ethernet synthetic loss measurement (ETH-SLM).

Single-ended ETH-SLM is used to perform on-demand or proactive tests by initiating a finite amount of ETH-SLM frames to one or multiple MEP peers and receiving the ETH-SLM reply from the peers. The ETH-SLM frames contain the ETH-SLM information that is used to measure and report both near-end and far-end synthetic loss measurements. Service-level agreement (SLA) measurement is the process of monitoring the bandwidth, delay, delay variation (jitter), continuity, and availability of a service. It enables you to identify network problems before customers are impacted by network defects. In proactive mode, SLA measurements are triggered by an iterator application. An iterator is designed to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for synthetic frame loss measurement. This mode differs from on-demand SLA measurement, which is user initiated. In on-demand mode, the measurements are triggered by the user through the CLI. When the user triggers the ETH-SLM through the CLI, the SLM request that is generated is as per the frame formats specified by the ITU-T Y.1731 standard.

**NOTE:** ACX5048 and ACX5096 routers support ETH-SLM for Layer 2 services.

SEE ALSO

[Troubleshooting Failures with ETH-SLM](#) | 809

## Scenarios for Configuration of ETH-SLM

ETH-SLM measures near-end and far-end frame loss between two MEPs that are part of the same MEG level. You can configure ETH-SLM to measure synthetic loss for both upward-facing or upstream MEP

and downward-facing or downstream MEP. This section describes the following scenarios for the operation of ETH-SLM:

#### ***Upstream MEP in MPLS Tunnels***

Consider a scenario in which a MEP is configured between the user network interfaces (UNIs) of two MX Series routers, MX1 and MX2, in the upstream direction. MX1 and MX2 are connected over an MPLS core network. ETH-SLM measurements are performed between the upstream MEP in the path linking the two routers. Both MX1 and MX2 can initiate on-demand or proactive ETH-SLM, which can measure both far-end and near-end loss at MX1 and MX2, respectively. The two UNIs are connected using MPLS-based Layer 2 VPN virtual private wire service (VPWS).

#### ***Downstream MEP in Ethernet Networks***

Consider a scenario in which a MEP is configured between two MX Series routers, MX1 and MX2, on the Ethernet interfaces in the downstream direction. MX1 and MX2 are connected in an Ethernet topology and downstream MEP is configured toward the Ethernet network. ETH-SLM measurements are performed between the downstream MEP in the path linking the two routers. ETH-SLM can be measured in the path between these two routers.

Consider another scenario in which a MEP is configured in the downstream direction and service protection for a VPWS over MPLS is enabled by specifying a working path or protect path on the MEP. Service protection provides end-to-end connection protection of the working path in the event of a failure. To configure service protection, you must create two separate transport paths—a working path and a protect path. You can specify the working path and protect path by creating two maintenance associations. To associate the maintenance association with a path, you must configure the MEP interface in the maintenance association and specify the path as working or protect.

In a sample topology, an MX Series router, MX1, is connected to two other MX Series routers, MX2 and MX3, over an MPLS core. The connectivity fault management (CFM) session between MX1 and MX2 is the working path on the MEP and the CFM session between MX1 and MX3 is the protect path on the MEP. MX2 and MX3 are, in turn, connected on Ethernet interfaces to MX4 in the access network. Downstream MEP is configured between MX1 and MX4 that passes through MX2 (working CFM session) and also between MX1 and MX4 that passes through MX3 (protected CFM session). ETH-SLM is performed between these downstream MEPs. In both the downstream MEPs, the configuration is performed on MX1 and MX4 UNIs, similar to upstream MEP.

SEE ALSO

| [Troubleshooting Failures with ETH-SLM](#) | 809

## Format of ETH-SLM Messages

Synthetic loss messages (SLMs) support single-ended Ethernet synthetic loss measurement (ETH-SLM) requests. This topic contains the following sections that describe the formats of the SLM protocol data units (PDUs), SLR PDUs, and the data iterator type length value (TLV).

### ***SLM PDU Format***

The SLM PDU format is used by a MEP to transmit SLM information. The following components are contained in SLM PDUs:

- **Source MEP ID**—Source MEP ID is a 2-octet field where the last 13 least significant bits are used to identify the MEP transmitting the SLM frame. MEP ID is unique within the MEG.
- **Test ID**—Test ID is a 4-octet field set by the transmitting MEP and is used to identify a test when multiple tests run simultaneously between MEPs (including both concurrent on-demand and proactive tests).
- **TxFCf**—TxFCf is a 4-octet field that carries the number of SLM frames transmitted by the MEP toward its peer MEP.

The following are the fields in an SLM PDU:

- **MEG Level**—Configured maintenance domain level in the range 0–7.
- **Version**—0.
- **OpCode**—Identifies an OAM PDU type. For SLM, it is 55.
- **Flags**—Set to all zeros.
- **TLV Offset**—16.
- **Source MEP ID**—A 2-octet field used to identify the MEP transmitting the SLM frame. In this 2-octet field, the last 13 least significant bits are used to identify the MEP transmitting the SLM frame. MEP ID is unique within the MEG.
- **RESV**—Reserved fields are set to all zeros.
- **Test ID**—A 4-octet field set by the transmitting MEP and used to identify a test when multiple tests run simultaneously between MEPs (including both concurrent on-demand and proactive tests).
- **TxFCf**—A 4-octet field that carries the number of SLM frames transmitted by the MEP toward its peer MEP.
- **Optional TLV**—A data TLV may be included in any SLM transmitted. For the purpose of ETH-SLM, the value part of data TLV is unspecified.
- **End TLV**—All zeros octet value.

### ***SLR PDU Format***

The synthetic loss reply (SLR) PDU format is used by a MEP to transmit SLR information. The following are the fields in an SLR PDU:

- MEG Level—A 3-bit field the value of which is copied from the last received SLM PDU.
- Version—A 5-bit field the value of which is copied from the last received SLM PDU.
- OpCode—Identifies an OAM PDU type. For SLR, it is set as 54.
- Flags—A 1-octet field copied from the SLM PDU.
- TLV Offset—A 1-octet field copied from the SLM PDU.
- Source MEP ID—A 2-octet field copied from the SLM PDU.
- Responder MEP ID—A 2-octet field used to identify the MEP transmitting the SLR frame.
- Test ID—A 4-octet field copied from the SLM PDU.
- TxFCf—A 4-octet field copied from the SLM PDU.
- TxFCb—A 4 octet field. This value represents the number of SLR frames transmitted for this test ID.
- Optional TLV—The value is copied from the SLM PDU, if present.
- End TLV—A 1-octet field copied from the SLM PDU.

#### ***Data Iterator TLV Format***

The data iterator TLV specifies the data TLV portion of the Y.1731 data frame. The MEP uses a data TLV when the MEP is configured to measure delay and delay variation for different frame sizes. The following are the fields in a data TLV:

- Type—Identifies the TLV type; value for this TLV type is Data (3).
- Length—Identifies the size, in octets, of the Value field containing the data pattern. The maximum value of the Length field is 1440.
- Data pattern—An  $n$ -octet ( $n$  denotes length) arbitrary bit pattern. The receiver ignores it.

SEE ALSO

[Troubleshooting Failures with ETH-SLM | 809](#)

### **Transmission of ETH-SLM Messages**

The ETH-SLM functionality can process multiple synthetic loss message (SLM) requests simultaneously between a pair of MEPs. The session can be a proactive or an on-demand SLM session. Each SLM request is identified uniquely by a test ID.

A MEP can send SLM requests or respond to SLM requests. A response to an SLM request is called a synthetic loss reply (SLR). After a MEP determines an SLM request by using the test ID, the MEP calculates the far-end and near-end frame loss on the basis of the information in the SLM message or the SLM protocol data unit (PDU).

A MEP maintains the following local counters for each test ID and for each peer MEP being monitored in a maintenance entity for which loss measurements are to be performed:

- **TxFCL**—Number of synthetic frames transmitted toward the peer MEP for a test ID. A source MEP increments this number for successive transmission of synthetic frames with ETH-SLM request information while a destination or receiving MEP increments this value for successive transmission of synthetic frames with the SLR information.
- **RxFCL**—Number of synthetic frames received from the peer MEP for a test ID. A source MEP increments this number for successive reception of synthetic frames with SLR information while a destination or receiving MEP increments it for successive reception of synthetic frames with ETH-SLM request information.

The following sections describe the phases of processing of SLM PDUs to determine synthetic frame loss:

#### ***Initiation and Transmission of SLM Requests***

A MEP periodically transmits an SLM request with the OpCode field set as 55. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TxFCL is sent in the packet.

No synchronization is required of the test ID value between initiating and responding MEPs because the test ID is configured at the initiating MEP, and the responding MEP uses the test ID it receives from the initiating MEP. Because ETH-SLM is a sampling technique, it is less precise than counting the service frames. Also, the accuracy of measurement depends on the number of SLM frames used or the period for transmitting SLM frames.

#### ***Reception of SLMs and Transmission of SLRs***

After the destination MEP receives a valid SLM frame from the source MEP, an SLR frame is generated and transmitted to the requesting or source MEP. The SLR frame is valid if the MEG level and the destination MAC address match the receiving MEP's MAC address. All the fields in the SLM PDUs are copied from the SLM request except for the following fields:

- The source MAC address is copied to the destination MAC address and the source address contains the MEP's MAC address.
- The value of the OpCode field is changed from SLM to SLR (54).
- The responder MEP ID is populated with the MEP's MEP ID.
- TxFCb is saved with the value of the local counter RxFCL at the time of SLR frame transmission.
- An SLR frame is generated every time an SLM frame is received; therefore, RxFCL in the responder is equal to the number of SLM frames received and also equal to the number of SLR frames sent. At the responder or receiving MEP, RxFCL equals TxFCL.

### **Reception of SLRs**

After an SLM frame (with a given TxFCf value) is transmitted, a MEP expects to receive a corresponding SLR frame (carrying the same TxTCf value) within the timeout value from its peer MEP. SLR frames that are received after the timeout value (5 seconds) are discarded. With the information contained in SLR frames, a MEP determines the frame loss for the specified measurement period. The measurement period is a time interval during which the number of SLM frames transmitted is statistically adequate to make a measurement at a given accuracy. A MEP uses the following values to determine near-end and far-end frame loss during the measurement period:

- Last received SLR frame's TxFCf and TxFCb values and the local counter RxFCI value at the end of the measurement period. These values are represented as TxFCf[tc], TxFCb[tc], and RxFCI[tc], where tc is the end time of the measurement period.
- SLR frame's TxFCf and TxFCb values of the first received SLR frame after the test starts and local counter RxFCI at the beginning of the measurement period. These values are represented as TxFCf[tp], TxFCb[tp], and RxFCI[tp], where tp is the start time of the measurement period.

For each SLR packet that is received, the local RxFCI counter is incremented at the sending or source MEP.

### **Computation of Frame Loss**

Synthetic frame loss is calculated at the end of the measurement period on the basis of the value of the local counters and the information from the last frame received. The last received frames contains the TxFCf and TxFCb values. The local counter contains the RxFCI value. Using these values, frame loss is determined using the following formula:

Frame loss (far-end) = TxFCf – TxFCb

Frame loss (near-end) = TxFCb – RxFCI

SEE ALSO

| [Troubleshooting Failures with ETH-SLM](#) | 809



## Release History Table

| Release              | Description                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">16.1</a> | Starting Junos OS Release 16.1, the Ethernet loss measurement (ETH-LM) results are inaccurate when connectivity fault management (CFM) and performance monitoring (PM) PDUs received locally at a maintenance endpoint (MEP) as classified as belonging to the yellow class or a packet loss priority (PLP) of medium-high.                                                                                    |
| <a href="#">16.1</a> | Starting with Junos OS Release 16.1, performance monitoring for connectivity fault management (by including the <b>performance-monitoring</b> statement and its substatements at the <b>[edit protocols oam ethernet connectivity-fault-management]</b> 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

[Configuring Ethernet Frame Loss Measurement | 747](#)

[Configuring an Iterator Profile | 781](#)

[Configuring Ethernet Synthetic Loss Measurements | 796](#)

## Configuring Ethernet Frame Delay Measurement Sessions

### IN THIS SECTION

- [Guidelines for Configuring Routers to Support an ETH-DM Session | 710](#)
- [Guidelines for Starting an ETH-DM Session | 711](#)
- [Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts | 713](#)
- [Configuring Routers to Support an ETH-DM Session | 718](#)
- [Triggering an Ethernet Frame Delay Measurements Session | 723](#)
- [Starting an ETH-DM Session | 725](#)
- [Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces | 728](#)
- [Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces | 734](#)
- [Managing Continuity Measurement Statistics | 740](#)
- [Viewing Ethernet Frame Delay Measurements Statistics | 741](#)
- [Managing ETH-DM Statistics and ETH-DM Frame Counts | 743](#)

Use this topic to understand how to configure Ethernet frame delay measurement sessions. You can start either a one-way Ethernet delay measurement session or a two-way Ethernet delay measurement session. Also, use this topic to view the delay measurement statistics and frame counts.

## Guidelines for Configuring Routers to Support an ETH-DM Session

### IN THIS SECTION

- [Configuration Requirements for ETH-DM | 710](#)
- [Configuration Options for ETH-DM | 711](#)

Keep the following guidelines in mind when configuring routers to support an Ethernet frame delay measurement (ETH-DM) session:

### **Configuration Requirements for ETH-DM**

You can obtain ETH-DM information for a link that meets the following requirements:

- The measurements can be performed between peer maintenance association endpoints (MEPs) on two routers.
- The two MEPs must be configured on two Ethernet physical interfaces or on two Ethernet logical interfaces. For more information, see [“Configuring a MEP to Generate and Respond to CFM Protocol Messages” on page 564](#).
- The two MEPs must be configured—on their respective routers—under the same maintenance association (MA) identifier. For more information, see [“Creating a Maintenance Association” on page 561](#).
- On both routers, the MA must be associated with the same maintenance domain (MD) name. For more information, see [“Creating a Maintenance Domain” on page 555](#).
- On both routers, periodic packet management (PPM) must be running on the Routing Engine and Packet Forwarding Engine, which is the default configuration. You can disable PPM on the Packet Forwarding Engine only. However, the Ethernet frame delay measurement feature requires that distributed PPM remain enabled on the Packet Forwarding Engine of both routers. For more information about **ppm**, see the *Junos OS Routing Protocols Library*.
- If the PPM process (**ppm**) is disabled on the Packet Forwarding Engine, you must re-enable it. Re-enabling distributed **ppm** entails restarting the **ethernet-connectivity-fault-management** process, which causes all connectivity fault management (CFM) sessions to re-establish. For more information about CFM sessions, see [“Configuring Ethernet Local Management Interface” on page 579](#).

**NOTE:** The Ethernet frame delay measurement feature is supported only for MEPs configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.

### **Configuration Options for ETH-DM**

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.

SEE ALSO

[Ethernet Alarm Indication | 810](#)

[Inline Transmission Mode | 824](#)

### **Guidelines for Starting an ETH-DM Session**

#### **IN THIS SECTION**

- [ETH-DM Session Prerequisites | 711](#)
- [ETH-DM Session Parameters | 712](#)
- [Restrictions for an ETH-DM Session | 713](#)

Keep the following guidelines in mind when preparing to start an Ethernet frame delay measurement (ETH-DM) session:

#### ***ETH-DM Session Prerequisites***

Before you can start an ETH-DM session, you must configure two MX Series routers to support ETH-DM by defining the two CFM-enabled physical or logical Ethernet interfaces on each router. This entails creating and configuring CFM maintenance domains, maintenance associations, and maintenance association end points on each router. For more information about enabling CFM on an Ethernet interface, see [“Creating a Maintenance Domain” on page 555](#).

**NOTE:** The Ethernet frame delay measurement feature is supported only for maintenance association end points configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.

For specific information about configuring routers to support ETH-DM, see [“Guidelines for Configuring Routers to Support an ETH-DM Session” on page 710](#) and [“Configuring Routers to Support an ETH-DM Session” on page 718](#).

### **ETH-DM Session Parameters**

You can initiate a one-way or two-way ETH-DM session by entering the **monitor ethernet delay-measurement** operational command at a router that contains one end of the service for which you want to measure frame delay. The command options specify the ETH-DM session in terms of the CFM elements:

- The type of ETH-DM measurement (one-way or two-way) to be performed.
- The Ethernet service for which the ETH-DM measurement is to be performed:
  - CFM maintenance domain—Name of the existing maintenance domain (MD) for which you want to measure Ethernet frame delays. For more information, see [“Creating a Maintenance Domain” on page 555](#).
  - CFM maintenance association—Name of an existing maintenance association (MA) within the maintenance domain. For more information, see [“Creating a Maintenance Association” on page 561](#).
  - Remote CFM maintenance association end point—The unicast MAC address or the numeric identifier of the remote maintenance association end point (MEP)—the physical or logical interface on the remote router that resides in the specified MD and is named in the specified MA—with which to perform the ETH-DM session. For more information, see [“Configuring a MEP to Generate and Respond to CFM Protocol Messages” on page 564](#).
- Optional specifications:
  - Count—You can specify the number of ETH-DM requests to send for this frame delay measurement session. The range is from 1 through 65,535 frames. The default value is 10 frames.
 

**NOTE:** Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).
  - Frame interval—You can specify the number of seconds to elapse between ETH-DM frame transmittals. The default value is 1 second.

For more detailed information about the parameters you can specify to start an ETH-DM session, see the **monitor ethernet delay-measurement** operational command description in the [CLI Explorer](#).

### Restrictions for an ETH-DM Session

The following restrictions apply to an ETH-DM session:

- You cannot run multiple simultaneous ETH-DM sessions with the same remote MEP or MAC address.
- For a given ETH-DM session, you can collect frame delay information for a maximum of 65,535 frames.
- For a given CFM session (pair of peer MEPs), the ETH-DM database stores a maximum of 100 statistics, with the older statistics being “aged out” as newer statistics are collected for that pair of MEPs.
  - For one-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the receiver MEP is defined.
  - For two-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the initiator MEP is defined.

Depending on the number of frames exchanged in the individual ETH-DM sessions, the ETH-DM database can contain statistics collected through multiple ETH-DM sessions.

- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *High Availability User Guide*.
- Accuracy of frame delay data is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

SEE ALSO

[Using the monitor ethernet delay-measurement Command | 725](#)  
[monitor ethernet delay-measurement | 1451](#)

### Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts

#### IN THIS SECTION

- [ETH-DM Statistics | 714](#)
- [ETH-DM Statistics Retrieval | 716](#)
- [ETH-DM Frame Counts | 716](#)
- [ETH-DM Frame Count Retrieval | 717](#)

### ETH-DM Statistics

Ethernet frame delay statistics are the frame delay and frame delay variation values determined by the exchange of frames containing ETH-DM protocol data units (PDUs).

- For a one-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, including the exchange of one-way delay PDU frames, see [“Ethernet Frame Delay Measurements Overview” on page 690](#).
- For a two-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, including the exchange of two-way delay PDU frames, see [“Ethernet Frame Delay Measurements Overview” on page 690](#).

A CFM database stores CFM-related statistics and—for Ethernet interfaces that support ETH-DM—the 100 most recently collected ETH-DM statistics for that pair of MEPs. You can view ETH-DM statistics by using the **delay-statistics** or **mep-statistics** form of the **show oam ethernet connectivity-fault-management** command to display the CFM statistics for the MEP that collects the ETH-DM statistics you want to view.

[Table 91 on page 714](#) describes the ETH-DM statistics calculated in an ETH-DM session.

**Table 91: ETH-DM Statistics**

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>One-way delay (μsec)<sup>†</sup></b> | <p>For a one-way ETH-DM session, the frame delay, in microseconds, collected at the receiver MEP.</p> <p>To display frame delay statistics for a given one-way ETH-DM session, use the <b>delay-statistics</b> or <b>mep-statistics</b> form of the <b>show oam ethernet connectivity-fault-management</b> command at the receiver MEP for that session.</p>                                                                                                                                                                                                                |
| <b>Two-way delay (μsec)</b>             | <p>For a two-way ETH-DM session, the frame delay, in microseconds, collected at the initiator MEP.</p> <p>When you start a two-way frame delay measurement, the CLI output displays each DMR frame receipt timestamp and corresponding DMM frame delay and delay variation collected as the session progresses.</p> <p>To display frame delay statistics for a given two-way ETH-DM session, use the <b>delay-statistics</b> or <b>mep-statistics</b> form of the <b>show oam ethernet connectivity-fault-management</b> command at the initiator MEP for that session.</p> |

Table 91: ETH-DM Statistics (*continued*)

| Field Name                                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Average delay</b> <sup>†</sup>           | <p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a <b>show</b> command, the <b>Average delay</b> field displays the average one-way and two- frame delays among all ETH-DM statistics collected at the CFM session level.</p> <p>For example, suppose you start two one-way ETH-DM sessions for 50 counts each, one after the other. If, after both measurement sessions complete, you use a <b>show</b> command to display 100 ETH-DM statistics for that CFM session, the <b>Average delay</b> field displays the average frame delay among all 100 statistics.</p> |
| <b>Average delay variation</b> <sup>†</sup> | <p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay variation among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a <b>show</b> command, the <b>Average delay variation</b> field displays the average one-way and two- frame delay variations among all ETH-DM statistics collected at the CFM session level.</p>                                                                                                                                                                                                                                                                                                           |
| <b>Best-case delay</b> <sup>†</sup>         | <p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the lowest two-way frame delay value among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a <b>show</b> command, the <b>Best case delay</b> field displays the lowest one-way and two-way frame delays among all ETH-DM statistics collected at the CFM session level.</p>                                                                                                                                                                                                                                                                                                                                |
| <b>Worst-case delay</b> <sup>†</sup>        | <p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the highest two-way frame delay value among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a <b>show</b> command, the <b>Worst case delay</b> field displays the highest one-way and two-way frame delays among all statistics collected at the CFM session level.</p>                                                                                                                                                                                                                                                                                                                                    |

<sup>†</sup>When you start a one-way frame delay measurement, the CLI output displays **NA** ("not available") for this field. One-way ETH-DM statistics are collected at the remote (receiver) MEP. Statistics for a given one-way ETH-DM session are available only by displaying CFM statistics for the receiver MEP.

### ETH-DM Statistics Retrieval

At the receiver MEP for a one-way session, or at the initiator MEP for a two-way session, you can display all ETH-DM statistics collected at a CFM session level by using the following operational commands:

- **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain *md-name* maintenance-association *ma-name* <local-mep *mep-id*> <remote-mep *mep-id*> <count *count*>**
- **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *md-name* maintenance-association *ma-name* <local-mep *mep-id*> <remote-mep *mep-id*> <count *count*>**

### ETH-DM Frame Counts

The number of ETH-DM PDU frames exchanged in a ETH-DM session are stored in the CFM database on each router.

[Table 92 on page 716](#) describes the ETH-DM frame counts collected in an ETH-DM session.

**Table 92: ETH-DM Frame Counts**

| Field Name                   | Field Description                                                                                                                                                                            |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>1DMs sent</b>             | <p>Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.</p> <p>Stored in the CFM database of the MEP initiating a one-way frame delay measurement.</p> |
| <b>Valid 1DMs received</b>   | <p>Number of valid 1DM frames received.</p> <p>Stored in the CFM database of the MEP receiving a one-way frame delay measurement.</p>                                                        |
| <b>Invalid 1DMs received</b> | <p>Number of invalid 1DM frames received.</p> <p>Stored in the CFM database of the MEP receiving a one-way frame delay measurement.</p>                                                      |
| <b>DMMs sent</b>             | <p>Number of delay measurement message (DMM) PDU frames sent to the peer MEP in this session.</p> <p>Stored in the CFM database of the MEP initiating a two-way frame delay measurement.</p> |
| <b>DMRs sent</b>             | <p>Number of delay measurement reply (DMR) frames sent (in response to a received DMM).</p> <p>Stored in the CFM database of the MEP responding to a two-way frame delay measurement.</p>    |
| <b>Valid DMRs received</b>   | <p>Number of valid DMR frames received.</p> <p>Stored in the CFM database of the MEP initiating a two-way frame delay measurement.</p>                                                       |
| <b>Invalid DMRs received</b> | <p>Number of invalid DMR frames received.</p> <p>Stored in the CFM database of the MEP initiating a two-way frame delay measurement.</p>                                                     |



### ***ETH-DM Frame Count Retrieval***

Each router counts the number of ETH-DM frames sent or received and stores the counts in a CFM database.

### ***Frame Counts Stored in CFM Databases***

You can display ETH-DM frame counts for MEPs assigned to specified Ethernet interfaces or for specified MEPs in CFM sessions by using the following operational commands:

- **show oam ethernet connectivity-fault-management interfaces (detail | extensive)**
- **show oam ethernet connectivity-fault-management mep-database maintenance-domain *md-name* maintenance-association *ma-name* <local-mep *mep-id*> <remote-mep *mep-id*>**

### ***One-Way ETH-DM Frame Counts***

For a one-way ETH-DM session, delay statistics are collected at the receiver MEP only, but frame counts are collected at both MEPs. As indicated in [Table 92 on page 716](#), one-way ETH-DM frame counts are tallied from the perspective of each router in the session:

- At the initiator MEP, the router counts the number of 1DM frames sent.
- At the receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received.

You can also view one-way ETH-DM frame counts—for a receiver MEP—by using the **show oam ethernet connectivity-fault-management mep-statistics** command to display one-way statistics and frame counts together.

### ***Two-Way ETH-DM Frame Counts***

For a two-way ETH-DM session, delay statistics are collected at the initiator MEP only, but frame counts are collected at both MEPs. As indicated in [Table 92 on page 716](#), two-way ETH-DM frame counts are tallied from the perspective of each router in the session:

- At the initiator MEP, the router counts the number of DMM frames sent, valid DMR frames received, and invalid DMR frames received.
- At the responder MEP, the router counts the number of DMR frames sent.

You can also view two-way ETH-DM frame counts—for an initiator MEP—by using the **show oam ethernet connectivity-fault-management mep-statistics** command to display two-way statistics and frame counts together.

SEE ALSO

---

[clear oam ethernet connectivity-fault-management statistics | 1444](#) command

---

[show oam ethernet connectivity-fault-management mep-statistics | 1916](#) command

---

[show oam ethernet connectivity-fault-management delay-statistics](#) | 1869 command

[show oam ethernet connectivity-fault-management interfaces](#) | 1879 (detail | extensive) command

[show oam ethernet connectivity-fault-management mep-database](#) | 1901 command

## Configuring Routers to Support an ETH-DM Session

### IN THIS SECTION

- [Configuring MEP Interfaces](#) | 718
- [Ensuring That Distributed ppm Is Not Disabled](#) | 719
- [Enabling the Hardware-Assisted Timestamping Option](#) | 722
- [Configuring the Server-Side Processing Option](#) | 722

### Configuring MEP Interfaces

Before you can start an Ethernet frame delay measurement session across an Ethernet service, you must configure two MX Series routers to support ETH-DM.

To configure an Ethernet interface on a MX Series router to support ETH-DM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

```
[edit interfaces]
interface {
  ethernet-interface-name {
    vlan-tagging;
    unit logical-unit-number {
      vlan-id vlan-id; # Both interfaces on this VLAN
    }
  }
}
```

Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

```
[edit protocols]
oam {
```

```

ethernet {
  connectivity-fault-management {
    maintenance-domain md-name { # On both routers
      level number;
      maintenance-association ma-name { # On both routers
        continuity-check {
          interval 100ms;
          hold-interval 1;
        }
        mep mep-id { # Attach to VLAN interface
          auto-discovery;
          direction (up | down);
          interface interface-name;
          priority number;
        }
      }
    }
  }
}

```

### ***Ensuring That Distributed ppm Is Not Disabled***

By default, the router's period packet management process (**ppm**) runs sessions distributed to the Packet Forwarding Engine in addition to the Routing Engine. This process is responsible for periodic transmission of packets on behalf of its various client processes, such as Bidirectional Forwarding Detection (BFD), and it also receives packets on behalf of client processes.

In addition, **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 such processes as BFD on the Packet Forwarding Engine.

### **Distributed ppm Required for ETH-DM**

Ethernet frame delay measurement requires that **ppm** remains distributed to the Packet Forwarding Engine. If **ppm** is not distributed to the Packet Forwarding Engines of both routers, ETH-DM PDU frame timestamps and ETH-DM statistics are not valid.

Before you start ETH-DM, you must verify that the following configuration statement is **NOT** present:

[edit]

```

routing-options {
  ppm {
    no-delegate-processing;
  }
}

```

If distributed **ppm** processing is disabled (as shown in the stanza above) on either router, you must re-enable it in order to use the ETH-DM feature.

### Procedure to Ensure that Distributed ppm is Not Disabled

To ensure that distributed **ppm** is not disabled on a router:

1. Display the packet processing management (PPM) configuration to determine whether distributed **ppm** is disabled.
  - In the following example, distributed **ppm** is enabled on the router. In this case, you do not need to modify the router configuration:

```

[edit]
user@host# show routing-options
ppm;

```

- In the following example, distributed **ppm** is disabled on the router. In this case, you must proceed to Step 2 to modify the router configuration:

```

[edit]
user@host# show routing-options
ppm {
  no-delegate-processing;
}

```

2. Modify the router configuration to re-enable distributed **ppm** and restart the Ethernet OAM Connectivity Fault Management process *ONLY IF* distributed **ppm** is disabled (as determined in the previous step).

- a. Before continuing, make any necessary preparations for the possible loss of connectivity on the router.

Restarting the **ethernet-connectivity-fault-management** process has the following effect on your network:

- All connectivity fault management (CFM) sessions re-establish.
- All ETH-DM requests on the router terminate.
- All ETH-DM statistics and frame counts reset to 0.

- b. Modify the router configuration to re-enable distributed **ppm**. For example:

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

- c. Commit the updated router configuration. For example:

```
[edit]
user@host# commit and-quit
commit complete
exiting configuration mode
```

- d. To restart the Ethernet OAM Connectivity-Fault-Management process, enter the **restart ethernet-connectivity-fault-management** **<gracefully | immediately | soft>** operational mode command. For example:

```
user@host> restart ethernet-connectivity-fault-management
Connectivity fault management process started, pid 9893
```

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.

SEE ALSO

[Periodic Packet Management | 183](#)

[Understanding Periodic Packet Management on MX Series Routers | 184](#)

### ***Enabling the Hardware-Assisted Timestamping Option***

By default, Ethernet frame delay measurement uses software for timestamping transmitted and received ETH-DM frames. For Ethernet interfaces, you can optionally use hardware timing to assist in the timestamping of received ETH-DM frames to increase the accuracy of delay measurements.

Enabling hardware-assisted timestamping of received frames can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction.

To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      performance-monitoring {
        hardware-assisted-timestamping;
      }
    }
  }
}
```

SEE ALSO

[hardware-assisted-timestamping | 1028](#)

### ***Configuring the Server-Side Processing Option***

You can delegate the server-side processing (for both two-way delay measurement and loss measurement) to the Packet Forwarding Engine to prevent overloading on the Routing Engine. By default, the server-side processing is done by the Routing Engine.

To configure the server-side processing option:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit protocols oam ethernet connectivity-fault-management performance-monitoring
```

2. Configure the server-side processing option.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set delegate-server-processing
```

3. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# show
performance-monitoring {
    delegate-server-processing;
}
```

#### SEE ALSO

[delegate-server-processing](#) | 1018

#### SEE ALSO

[Understanding Periodic Packet Management on MX Series Routers](#) | 184  
[Inline Transmission Mode](#) | 824

### Triggering an Ethernet Frame Delay Measurements Session

Before Ethernet frame delay measurement statistics can be displayed, they must be collected. To trigger Ethernet frame delay measurement, use the **monitor ethernet delay-measurement (one-way | two-way) (remote-mac-address | mep identifier) maintenance-domain name maintenance-association ma-id [count count] [wait time]** operational command.

The fields for this command are described in [Table 93 on page 723](#).

**Table 93: Monitor Ethernet Delay Command Parameters**

| Parameter                        | Parameter Range     | Description                                                                                                                                         |
|----------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>one-way</b> or <b>two-way</b> | NA                  | Perform a one-way or two-way (round-trip) delay measurement.                                                                                        |
| <b>remote-mac-address</b>        | Unicast MAC address | Send delay measurement frames to the destination unicast MAC address (use the format xx:xx:xx:xx:xx:xx). Multicast MAC addresses are not supported. |

Table 93: Monitor Ethernet Delay Command Parameters (*continued*)

| Parameter                            | Parameter Range            | Description                                                                                                |
|--------------------------------------|----------------------------|------------------------------------------------------------------------------------------------------------|
| <b>mep identifier</b>                | 1-8191                     | The MEP identifier to use for the measurement. The discovered MAC address for this MEP identifier is used. |
| <b>maintenance-domain name</b>       | Existing MD name           | Specifies an existing maintenance domain (MD) to use for the measurement.                                  |
| <b>maintenance-association ma-id</b> | Existing MA identifier     | Specifies an existing maintenance association (MA) identifier to use for the measurement.                  |
| <b>count count</b>                   | 1-65535 (default: 10)      | (Optional) Specifies the number of Ethernet frame delay frames to send. The default is 10.                 |
| <b>wait time</b>                     | 1-255 seconds (default: 1) | (Optional) Specifies the number of seconds to wait between frames. The default is 1 second.                |

If you attempt to monitor delays to a nonexistent MAC address, you must exit the application manually using **^C**:

```
user@host> monitor ethernet delay-measurement two-way 00:11:22:33:44:55
```

```
Two-way ETH-DM request to 00:11:22:33:44:55, Interface ge-5/2/9.0
^C
--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 0
Average delay: 0 usec, Average delay variation: 0 usec
Best case delay: 0 usec, Worst case delay: 0 usec
```

SEE ALSO

| [Configuring Ethernet Frame Loss Measurement](#) | 747



## Starting an ETH-DM Session

### IN THIS SECTION

- [Using the monitor ethernet delay-measurement Command | 725](#)
- [Starting a One-Way ETH-DM Session | 726](#)
- [Starting a Two-Way ETH-DM Session | 727](#)

### *Using the monitor ethernet delay-measurement Command*

After you have configured two MX Series routers to support ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM), you can initiate a one-way or two-way Ethernet frame delay measurement session from the CFM maintenance association end point (MEP) on one of the routers to the peer MEP on the other router.

To start an ETH-DM session between the specified local MEP and the specified remote MEP, enter the **monitor ethernet delay-measurement** command at operational mode. The syntax of the command is as follows:

```
monitor ethernet delay-measurement
(one-way | two-way)
maintenance-domain md-name
maintenance-association ma-name
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<wait interval-seconds>
<priority 802.1p value>
<size>
<no-session-id-tlv>
<xml>
```

For a one-way frame delay measurement, the command displays a runtime display of the number of 1DM frames sent from the initiator MEP during that ETH-DM session. One-way frame delay and frame delay variation measurements from an ETH-DM session are collected in a CFM database at the router that contains the receiver MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

For a two-way frame delay measurement, the command displays two-way frame delay and frame delay variation values for each round-trip frame exchange during that ETH-DM session, as well as a runtime display of useful summary information about the session: average delay, average delay variation, best-case delay, and worst-case delay. Two-way frame delay and frame delay variation values measurements from

an ETH-DM session are collected in a CFM database at the router that contains the initiator MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

**NOTE:** Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).

For a complete description of the **monitor ethernet delay-measurement** operational command, see the [CLI Explorer](#).

SEE ALSO

| [monitor ethernet delay-measurement](#) | 1451

#### ***Starting a One-Way ETH-DM Session***

To start a one-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement one-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```

```
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

**NOTE:** If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

SEE ALSO

[monitor ethernet delay-measurement](#) | 1451

### *Starting a Two-Way ETH-DM Session*

To start a two-way Ethernet frame delay measurement session, enter the **monitor ethernet delay-measurement two-way** command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```

```
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec
```

**NOTE:** If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

SEE ALSO

[monitor ethernet delay-measurement](#) | 1451

SEE ALSO

| [Inline Transmission Mode](#) | 824

### Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces

This example uses two MX Series routers: **MX-1** and **MX-2**. The configuration creates a CFM down MEP session on a VLAN-tagged logical interface connecting the two (**ge-5/2/9** on Router **MX-1** and **ge-0/2/5** on Router **MX-2**).

**NOTE:** These are not complete router configurations.

#### Configuration on Router **MX-1**:

```
[edit]
interfaces {
  ge-5/2/9 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        linktrace {
          path-database-size 255;
          age 10s;
        }
        maintenance-domain md6 {
          level 6;
          maintenance-association ma6 {
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
          }
        }
      }
    }
  }
}
```



```

        interface ge-0/2/5.0;
        direction down;
        auto-discovery;
    }
}
}
}
}
}
}
}

```

From Router **MX-2**, start a one-way delay measurement to Router **MX-1**.

```

user@MX-2> monitor ethernet delay-measurement one-way mep 201 maintenance-domain md6
maintenance-association ma6 count 10

```

```

One-way ETH-DM request to 00:90:69:0a:43:94, Interface ge-0/2/5.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA

```

The counters are displayed as part of the local MEP database on Router **MX-2**.

```

user@MX-2> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md6 maintenance-domain ma6

```

```

Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : no
  Cross-connect CCM received                   : no
  RDI sent by some MEP                        : no
Statistics:
  CCMs sent                                   : 1590
  CCMs received out of sequence                : 0

```

```

LBMs sent : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
lDMs sent : 10
Valid lDMs received : 0
Invalid lDMs received : 0
DMMs sent : 0
DMRs sent : 0
Valid DMRs received : 0
Invalid DMRs received : 0
Remote MEP count: 1
  Identifier   MAC address   State   Interface
    201       00:90:69:0a:43:94   ok    ge-0/2/5.0

```

The remote MEP database statistics are available on Router **MX-1**.

user@MX-1> **show oam ethernet connectivity-fault-management mep-database maintenance-domain md6**

```

Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 201, Direction: down, MAC address: 00:90:69:0a:43:94
Auto-discovery: enabled, Priority: 0
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM : no
  Erroneous CCM received : no
  Cross-connect CCM received : no
  RDI sent by some MEP : no
Statistics:
  CCMs sent : 1572
  CCMs received out of sequence : 0
  LBMs sent : 0
  Valid in-order LBRs received : 0
  Valid out-of-order LBRs received : 0

```

```

LBRs received with corrupted data      : 0
LBRs sent                              : 0
LTMs sent                              : 0
LTMs received                          : 0
LTRs sent                              : 0
LTRs received                          : 0
Sequence number of next LTM request    : 0
lDMs sent                              : 0
Valid lDMs received                    : 10
Invalid lDMs received                  : 0
DMMs sent                              : 0
DMRs sent                              : 0
Valid DMRs received                   : 0
Invalid DMRs received                  : 0
Remote MEP count: 1
Identifier    MAC address      State    Interface
  101        00:90:69:0a:48:57   ok      ge-5/2/9.0

```

The remote Router **MX-1** should also collect the delay statistics (up to 100 per session) for display with **mep-statistics** or **delay-statistics**.

user@MX-1> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md6**

```

MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
  CCMs sent                              : 3240
  CCMs received out of sequence          : 0
  LBMs sent                              : 0
  Valid in-order LBRs received           : 0
  Valid out-of-order LBRs received       : 0
  LBRs received with corrupted data      : 0
  LBRs sent                              : 0
  LTMs sent                              : 0
  LTMs received                          : 0
  LTRs sent                              : 0
  LTRs received                          : 0
  Sequence number of next LTM request    : 0
  lDMs sent                              : 0
  Valid lDMs received                    : 10
  Invalid lDMs received                  : 0
  DMMs sent                              : 0
  DMRs sent                              : 0

```



```
Valid DMRs received           : 0
Invalid DMRs received         : 0
```

```
Remote MEP identifier: 101
```

```
Remote MAC address: 00:90:69:0a:48:57
```

```
Delay measurement statistics:
```

```
Index  One-way delay  Two-way delay
      (usec)         (usec)
```

```
1      370
```

```
2      357
```

```
3      344
```

```
4      332
```

```
5      319
```

```
6      306
```

```
7      294
```

```
8      281
```

```
9      269
```

```
10     255
```

```
Average one-way delay       : 312 usec
```

```
Average one-way delay variation: 11 usec
```

```
Best case one-way delay     : 255 usec
```

```
Worst case one-way delay    : 370 usec
```

```
user@MX-1> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain
md6
```

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
```

```
Remote MEP count: 1
```

```
Remote MAC address: 00:90:69:0a:48:57
```

```
Delay measurement statistics:
```

```
Index  One-way delay  Two-way delay
      (usec)         (usec)
```

```
1      370
```

```
2      357
```

```
3      344
```

```
4      332
```

```
5      319
```

```
6      306
```

```
7      294
```

```
8      281
```

```
9      269
```

```
10     255
```

```

Average one-way delay      : 312 usec
Average one-way delay variation: 11 usec
Best case one-way delay    : 255 usec

```

**NOTE:** When two systems are close to each other, their one-way delay values are very high compared to their two-way delay values. This is because one-way delay measurement requires the timing for the two systems to be synchronized at a very granular level and MX Series routers do not support this granular synchronization. However, two-way delay measurement does not require synchronized timing, making two-way delay measurements more accurate.

## SEE ALSO

*Ethernet Interfaces User Guide for Routing Devices*

[Ethernet Frame Delay Measurements Overview | 690](#)

[Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements | 568](#)

[Triggering an Ethernet Frame Delay Measurements Session | 723](#)

[Viewing Ethernet Frame Delay Measurements Statistics | 741](#)

[Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces | 734](#)

*Configuring ETH-DM with Untagged Interfaces*

## Example: Configuring Two-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces

This example uses two MX Series routers: **MX-1** and **MX-2**. The configuration creates a CFM down MEP session on a VLAN-tagged logical interface connecting the two (**ge-5/2/9** on Router **MX-1** and **ge-0/2/5** on Router **MX-2**).

**NOTE:** These are not complete router configurations.

### Configuration on Router **MX-1**:

[edit]

```

interfaces {
  ge-5/2/9 {
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        linktrace {
          path-database-size 255;
          age 10s;
        }
        maintenance-domain md6 {
          level 6;
          maintenance-association ma6 {
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 201 {
              interface ge-5/2/9.0;
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}

```

Configuration on Router **MX-2**:

```

[edit]
interfaces {
  ge-0/2/5 {

```

```

    vlan-tagging;
    unit 0 {
        vlan-id 512;
    }
}
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                traceoptions {
                    file eoam_cfm.log size 1g files 2 world-readable;
                    flag all;
                }
                linktrace {
                    path-database-size 255;
                    age 10s;
                }
                maintenance-domain md6 {
                    level 6;
                    maintenance-association ma6 {
                        continuity-check {
                            interval 100ms;
                            hold-interval 1;
                        }
                        mep 101 {
                            interface ge-0/2/5.0;
                            direction down;
                            auto-discovery;
                        }
                    }
                }
            }
        }
    }
}
}
}

```

From Router **MX-1**, start a two-way delay measurement to Router **MX-2**.

```

user@MX-1> monitor ethernet delay-measurement two-way mep 101 maintenance-domain md6
maintenance-association ma6 count 10

```

```

Two-way ETH-DM request to 00:90:69:0a:48:57, Interface ge-5/2/9.0
DMR received from 00:90:69:0a:48:57 Delay: 100 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 8 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 111 usec Delay variation: 19 usec
DMR received from 00:90:69:0a:48:57 Delay: 110 usec Delay variation: 1 usec
DMR received from 00:90:69:0a:48:57 Delay: 119 usec Delay variation: 9 usec
DMR received from 00:90:69:0a:48:57 Delay: 122 usec Delay variation: 3 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 30 usec
DMR received from 00:90:69:0a:48:57 Delay: 92 usec Delay variation: 0 usec
DMR received from 00:90:69:0a:48:57 Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec

```

The counters are displayed as part of the MEP database on Router **MX-1** maintenance domain **MD6**.

**user@MX-1> show oam ethernet connectivity-fault-management mep-database maintenance-domain md6**

```

Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 201, Direction: down, MAC address: 00:90:69:0a:43:94
Auto-discovery: enabled, Priority: 0
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMS sent                                  : 894
  CCMS received out of sequence              : 0
  LBMS sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMS sent                                  : 0
  LTMS received                              : 0

```

```

LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
lDMs sent : 0
Valid lDMs received : 0
Invalid lDMs received : 0
DMMs sent : 10
DMRs sent : 0
Valid DMRs received : 10
Invalid DMRs received : 0
Remote MEP count: 1
Identifier    MAC address    State    Interface
101          00:90:69:0a:48:57    ok      ge-5/2/9.0

```

The collected MEP statistics are saved (up to 100 per remote MEP or per CFM session) and displayed as part of the MEP statistics on Router **MX-1**.

**user@MX-1> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md6**

```

MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
CCMs sent : 3154
CCMs received out of sequence : 0
LBMs sent : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
lDMs sent : 0
Valid lDMs received : 0
Invalid lDMs received : 0
DMMs sent : 10
DMRs sent : 0
Valid DMRs received : 10
Invalid DMRs received : 0

Remote MEP identifier: 101

```

```

Remote MAC address: 00:90:69:0a:48:57
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)         (usec)
  1             100
  2              92
  3              92
  4             111
  5             110
  6             119
  7             122
  8              92
  9              92
 10             108
Average two-way delay      : 103 usec
Average two-way delay variation: 8 usec
Best case two-way delay    : 92 usec
Worst case two-way delay   : 122 usec

```

The collected delay statistics are also saved (up to 100 per session) and displayed as part of the MEP delay statistics on Router **MX-1**.

user@MX-1> **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md6**

```

MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1

Remote MAC address: 00:90:69:0a:48:57
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)         (usec)
  1             100
  2              92
  3              92
  4             111
  5             110
  6             119
  7             122
  8              92
  9              92
 10             108
Average two-way delay      : 103 usec

```

```

Average two-way delay variation: 8 usec
Best case two-way delay       : 92 usec
Worst case two-way delay      : 122 usec

```

## SEE ALSO

*Ethernet Interfaces User Guide for Routing Devices*

[Ethernet Frame Delay Measurements Overview | 690](#)

[Configuring MEP Interfaces to Support Ethernet Frame Delay Measurements | 568](#)

[Triggering an Ethernet Frame Delay Measurements Session | 723](#)

[Viewing Ethernet Frame Delay Measurements Statistics | 741](#)

[Example: Configuring One-Way Ethernet Frame Delay Measurements with Single-Tagged Interfaces | 728](#)

*Configuring ETH-DM with Untagged Interfaces*

## Managing Continuity Measurement Statistics

### IN THIS SECTION

- [Displaying Continuity Measurement Statistics | 740](#)
- [Clearing Continuity Measurement Statistics | 741](#)

### *Displaying Continuity Measurement Statistics*

#### Purpose

Display continuity measurement.

The **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1** command is enhanced to display continuity measurement statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

#### Action

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

```

user@host> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain
md1 maintenance-association ma1

```



SEE ALSO

| [show oam ethernet connectivity-fault-management delay-statistics](#) | 1869

### **Clearing Continuity Measurement Statistics**

#### **Purpose**

Clear the continuity measurement statistics

By default, statistics are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

#### **Action**

- To clear the continuity measurement statistics for all MEPs attached to CFM-enabled interfaces on the router:

```
user@host> clear oam ethernet connectivity-fault-management continuity-measurement
maintenance-domain md-name maintenance-association ma-name local-mep local-mep-id remote-mep
remote-mep-id
```

SEE ALSO

| [clear oam ethernet connectivity-fault-management continuity-measurement](#) | 1437

SEE ALSO

| [clear oam ethernet connectivity-fault-management continuity-measurement](#) | 1437  
| [show oam ethernet connectivity-fault-management delay-statistics](#) | 1869

### **Viewing Ethernet Frame Delay Measurements Statistics**

Once Ethernet frame delay measurement statistics have been collected, they can be displayed.

To retrieve the last 100 Ethernet frame delay measurement statistics per remote MEP or per CFM session, two types of **show** commands are provided:

- For all OAM frame counters and Ethernet frame delay measurement statistics
- For Ethernet frame delay measurement statistics only

To retrieve all Ethernet frame delay measurement statistics for a given session, use the **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain *name* maintenance-association *name* [local-mep *identifier*] [remote-mep *identifier*] [count *count*]** command.

To retrieve only Ethernet frame delay measurement statistics for a given session, use the **show oam ethernet connectivity-fault-management delay-statistics maintenance-domain *name* maintenance-association *name* [local-mep *identifier*] [remote-mep *identifier*] [count *count*]** command.

**NOTE:** The only difference in the two commands is the use of the **mep-statistics** and **delay-statistics** keyword.

The fields for these commands are described in [Table 94 on page 742](#).

**Table 94: Show Ethernet Delay Command Parameters**

| Parameter                                   | Parameter Range        | Description                                                                                                     |
|---------------------------------------------|------------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>maintenance-domain <i>name</i></b>       | Existing MD name       | Specifies an existing maintenance domain (MD) to use.                                                           |
| <b>maintenance-association <i>ma-id</i></b> | Existing MA identifier | Specifies an existing maintenance association (MA) identifier to use.                                           |
| <b>local-mep <i>identifier</i></b>          | 1-8191                 | When a MEP has been specified, display statistics only for the local MEP.                                       |
| <b>remote-mep <i>identifier</i></b>         | 1-8191                 | When a MEP has been specified, display statistics only for the discovered MEP.                                  |
| <b>count <i>count</i></b>                   | 1-100 (default:100)    | The number of entries to display in the results table. By default, all 100 entries are displayed if they exist. |

**NOTE:** For each MEP, you will see frame counters for sent and received Ethernet frame delay measurement frames whenever MEP statistics are displayed.

SEE ALSO

| *Configuring MEPs to Generate and Respond to CFM Protocol Messages in a I*

## Managing ETH-DM Statistics and ETH-DM Frame Counts

### IN THIS SECTION

- [Displaying ETH-DM Statistics Only | 743](#)
- [Displaying ETH-DM Statistics and Frame Counts | 744](#)
- [Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity | 744](#)
- [Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level | 745](#)
- [Clearing ETH-DM Statistics and Frame Counts | 746](#)

### *Displaying ETH-DM Statistics Only*

#### Purpose

Display ETH-DM statistics.

By default, the **show oam ethernet connectivity-fault-management delay-statistics** command displays ETH-DM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

#### Action

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain  
ma1 maintenance-association ma1
```

- To display the ETH-DM statistics collected for ETH-DM sessions for the local MEP **201** belonging to MA **ma2** and within MD **md2**:

```
user@host> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain  
md2 maintenance-association ma2 local-mep 201
```

- To display the ETH-DM statistics collected for ETH-DM sessions from local MEPs belonging to MA **ma3** and within MD **md3** to remote MEP **302**:

```
user@host> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain  
md3 maintenance-association ma3 remote-mep 302
```

### SEE ALSO

| [show oam ethernet connectivity-fault-management delay-statistics](#) | 1869

### ***Displaying ETH-DM Statistics and Frame Counts***

#### **Purpose**

Display ETH-DM statistics and ETH-DM frame counts.

By default, the **show oam ethernet connectivity-fault-management mep-statistics** command displays ETH-DM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

#### **Action**

- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA **ma1** and within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain
md1 maintenance-association ma1
```

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP **201** in MA **ma2** and within MD **md2**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain
md2 maintenance-association ma2 local-mep 201
```

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP in MD **md3** and within MA **ma3** that participates in an ETH-DM session with the remote MEP **302**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain
ma3 maintenance-association ma3 remote-mep 302
```

SEE ALSO

| [show oam ethernet connectivity-fault-management mep-statistics](#) | 1916

### ***Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity***

#### **Purpose**

Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management mep-database** command displays CFM database information for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

**NOTE:** At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

### Action

- To display CFM database information (including ETH-DM frame counts) for all MEPs in MA **ma1** within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
ma1 maintenance-association ma1
```

- To display CFM database information (including ETH-DM frame counts) only for local MEP **201** in MA **ma1** within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md2 maintenance-association ma2 local-mep 201
```

- To display CFM database information (including ETH-DM frame counts) only for remote MEP **302** in MD **md3** within MA **ma3**:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
ma3 maintenance-association ma3 remote-mep 302
```

### SEE ALSO

[show oam ethernet connectivity-fault-management mep-database](#) | 1901

### *Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level*

#### Purpose

Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management interfaces** command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-DM, any frame counts are also displayed when you specify the **detail** or **extensive** command option.

**NOTE:** At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

**Action**

- To display CFM database information (including ETH-DM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
```

- To display CFM database information (including ETH-DM frame counts) only for the MEPs attached to CFM-enabled router interface **ge-5/2/9.0**:

```
user@host> show oam ethernet connectivity-fault-management interfaces ge-5/2/9.0 detail
```

- To display CFM database information (including ETH-DM frame counts) only for MEPs enclosed within CFM maintenance domains (MDs) at level **6**:

```
user@host> show oam ethernet connectivity-fault-management interfaces level 6 detail
```

## SEE ALSO

| [show oam ethernet connectivity-fault-management interfaces | 1879](#)

***Clearing ETH-DM Statistics and Frame Counts*****Purpose**

Clear the ETH-DM statistics and ETH-DM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

**Action**

- To clear the ETH-DM statistics and ETH-DM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

```
user@host> clear oam ethernet connectivity-fault-management statistics
```

- To clear the ETH-DM statistics and ETH-DM frame counts only for MEPs attached to the logical interface **ge-0/5.9.0**:

```
user@host> clear oam ethernet connectivity-fault-management statistics ge-0/5/9.0
```

## SEE ALSO

| [clear oam ethernet connectivity-fault-management statistics | 1444](#)

## SEE ALSO

[clear oam ethernet connectivity-fault-management statistics | 1444](#)[show oam ethernet connectivity-fault-management delay-statistics | 1869](#)[show oam ethernet connectivity-fault-management interfaces | 1879](#)[show oam ethernet connectivity-fault-management mep-statistics | 1916](#)[show oam ethernet connectivity-fault-management mep-database | 1901](#)

## RELATED DOCUMENTATION

[Configuring an Iterator Profile | 781](#)[Configuring Ethernet Synthetic Loss Measurements | 796](#)[Ethernet Alarm Indication | 810](#)[Inline Transmission Mode | 824](#)

## Configuring Ethernet Frame Loss Measurement

## IN THIS SECTION

- [Configuring Statistical Frame Loss Measurement for VPLS Connections | 747](#)
- [Managing ETH-LM Statistics | 748](#)
- [Example: Measuring Ethernet Frame Loss for Single-Tagged LMM/LMR PDUs | 750](#)
- [Example: Measuring Ethernet Frame Loss for Dual-Tagged LMM/LMR PDUs | 766](#)

Use this topic to understand more about frame loss measurement and how to configure frame loss measurement.

### Configuring Statistical Frame Loss Measurement for VPLS Connections

Using proactive statistical frame loss measurement, you can monitor VPLS connections on MX Series routers. Statistical frame loss measurement allows you to monitor the quality of Ethernet connections for service level agreements (SLAs). Point-to-point and multipoint-to-multipoint connections configured on MX Series routers can be monitored by registering the connection on an iterator and initiating periodic SLA measurement of frame transmissions on the connections.

Iterators periodically transmit SLA measurement packets using ITU-Y.1731 compliant frames. The iterator sends periodic measurement packets for each of the connections registered to it. These measurement cycles are transmitted in such a way as to not overlap, reducing the processing demands placed on the CPU. The measurement packets are exchanged between the source user network interface (UNI) port and the destination UNI port, providing a sequence of timed performance measurements for each UNI pair. The Frame Loss Ratio (FLR) and connection availability can be computed from these measurements using statistics.

The following steps outline how to configure statistical frame loss measurement for VPLS connections:

1. To configure proactive ETH-DM measurement for a VPLS connection, see [“Guidelines for Configuring Routers to Support an ETH-DM Session” on page 710](#).
2. To enable statistical loss measurement for a VPLS connection, configure an iterator for the VPLS connection using the [sla-iterator-profiles](#) statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level. For detailed instructions, see [“Configuring an Iterator Profile” on page 782](#).
3. As part of the iterator configuration, include the **statistical-frame-loss** option for the [measurement-type](#) statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles *profile-name*]** hierarchy level.
4. Once you have enabled the iterator, you can display the statistical frame loss for a VPLS connection by issuing the **show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator *identifier* maintenance-domain *name* maintenance-association *name* local-mep *identifier* remote-mep *identifier*** command.

SEE ALSO

[Configuring an Iterator Profile | 782](#)

[Verifying the Configuration of an Iterator Profile | 785](#)

## Managing ETH-LM Statistics

### IN THIS SECTION

- [Displaying ETH-LM Statistics | 749](#)
- [Clearing ETH-LM Statistics | 750](#)



## Displaying ETH-LM Statistics

### Purpose

Display the ETH-LM statistics.

By default, the **show oam ethernet connectivity-fault-management loss-statistics maintenance-domain *md-name* maintenance-association *ma-name*** command displays ETH-LM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

The following list consists of the CFM-related operational mode commands that have been enhanced to display ETH-LM statistics:

- The **show oam ethernet connectivity-fault-management interfaces detail** command is enhanced to display ETH-DM and ETH-LM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).
- The **show oam ethernet connectivity-fault-management mep-statistics** command is enhanced to display ETH-DM and ETH-LM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).
- The **show oam ethernet connectivity-fault-management mep-database** command is enhanced to display ETH-DM and ETH-LM frame counters for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

### Action

- To display the ETH-LM statistics for all MEPs attached to CFM-enabled interfaces on the router:

```
user@host> show oam ethernet connectivity-fault-management loss-statistics
```

- To display the ETH-DM statistics collected for MEPs belonging to MA **ma1** and within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md1 maintenance-association ma1
```

- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA **ma1** and within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1
```

- To display CFM database information (including ETH-DM frame counts) for all MEPs in MA **ma1** within MD **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md1 maintenance-association ma1
```

## SEE ALSO

|                                                                                                        |
|--------------------------------------------------------------------------------------------------------|
| <a href="#">clear oam ethernet connectivity-fault-management loss-statistics   1440</a>                |
| <a href="#">show oam ethernet connectivity-fault-management delay-statistics   1869</a>                |
| <a href="#">show oam ethernet connectivity-fault-management interfaces   1879 (detail   extensive)</a> |
| <a href="#">show oam ethernet connectivity-fault-management mep-statistics   1916</a>                  |
| <a href="#">show oam ethernet connectivity-fault-management mep-database   1901</a>                    |
| <a href="#">show oam ethernet connectivity-fault-management loss-statistics   1896</a>                 |

**Clearing ETH-LM Statistics****Purpose**

Clear the ETH-LM statistics.

By default, statistics are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

**Action**

- To clear the ETH-LM statistics for all MEPs attached to CFM-enabled interfaces on the router:

```
user@host> clear oam ethernet connectivity-fault-management loss-statistics
```

## SEE ALSO

|                                                                                         |
|-----------------------------------------------------------------------------------------|
| <a href="#">clear oam ethernet connectivity-fault-management loss-statistics   1440</a> |
|-----------------------------------------------------------------------------------------|

## SEE ALSO

|                                                                          |
|--------------------------------------------------------------------------|
| <a href="#">Managing ETH-DM Statistics and ETH-DM Frame Counts   743</a> |
|--------------------------------------------------------------------------|

**Example: Measuring Ethernet Frame Loss for Single-Tagged LMM/LMR PDUs****IN THIS SECTION**

- [Requirements | 751](#)
- [Overview and Topology | 751](#)
- [Configuration | 752](#)
- [Verification | 763](#)

This example illustrates how to configure Ethernet frame loss measurement (ETH-LM) for single-tagged Loss Measurement Message (LMM)/Loss Measurement Reply (LMR) protocol data units (PDUs). By configuring ETH-LM, you can measure the Ethernet frame loss that occur in your network.

### Requirements

This example uses the following hardware and software components:


- Two MX Series 5G Universal Routing Platforms with Rev-B Dense Port Concentrators (DPCs)
- Junos OS Release 14.2 or later

### Overview and Topology

Junos OS supports Ethernet frame loss measurement (ETH-LM) between maintenance association end points (MEPs) configured on Ethernet physical or logical interfaces on Rev-B Dense Port Concentrators (DPCs) in MX Series routers. Additionally, the Y.1731 functionality supports ETH-LM only for an end-to-end connection that uses Virtual Private Wire Service (VPWS). This example illustrates how to configure ETH-LM for single-tagged LMM/LMR PDUs with input and output VLAN map configured as **swap**.

[Figure 25 on page 751](#) shows the topology used in this example. VPWS service is configured between two MX Series routers, MX-PE1 and MX-PE2.

**Figure 25: VPWS Service Configured Between Two MX Series Routers**



MX-PE1 router has two Ethernet interfaces, **ge-5/0/4** and **ge-5/1/9**. Virtual LAN (VLAN) is configured on **ge-5/0/4** and MPLS is configured on the **ge-5/1/9** interface. The **ge-5/0/4.11** interface is used to configure the Layer 2 virtual circuit with MX-PE2 router. The UP MEP, **mep 2**, is attached to the **ge-5/0/4.11** interface. The three-color policer firewall filter is also configured for the MX-PE1 router.

Similarly, MX-PE2 router has two Ethernet interfaces, **ge-8/0/8** and **ge-8/0/9**. Virtual LAN (VLAN) is configured on **ge-8/0/8** and MPLS is configured on the **ge-8/0/9** interface. The **ge-8/0/8.11** interface is used to configure the Layer 2 virtual circuit with MX-PE1 router. The UP MEP, **mep 1**, is attached to the **ge-8/0/8.11** interface. The three-color policer firewall filter is also configured for the MX-PE2 router.

## Configuration

### IN THIS SECTION

- [Configuring Router PE1 | 754](#)
- [Configuring Router PE2 | 759](#)

### CLI Quick Configuration

To quickly configure ETH-LM for single-tagged LMM/LMR PDUs, copy the following commands, remove any line breaks, and then paste the commands into the CLI of each device.

On Router PE1:

```
[edit]
set interfaces ge-5/0/4 encapsulation flexible-ethernet-services
set interfaces ge-5/0/4 unit 11 encapsulation vlan-ccc
set interfaces ge-5/0/4 unit 11 layer2-policer input-three-color abc
set interfaces ge-5/0/4 unit 11 family ccc
set interfaces ge-5/1/9 enable
set interfaces ge-5/1/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-5/1/9 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 4.4.4.4/32
set interfaces ge-5/0/4 flexible-vlan-tagging
set interfaces ge-5/0/4 unit 11 vlan-id 2000
set interfaces ge-5/0/4 unit 11 input-vlan-map swap
set interfaces ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-5/0/4 unit 11 output-vlan-map swap
set routing-options router-id 4.4.4.4
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word
set protocols oam ethernet connectivity-fault-management performance-monitoring delegate-server-processing
set protocols oam ethernet connectivity-fault-management maintenance-domain md level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 2 interface ge-5/0/4.11
```

```

set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 2 direction up
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 2 remote-mep 1
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k

```

On Router PE2:

```

[edit]
set interfaces ge-8/0/8 encapsulation flexible-ethernet-services
set interfaces ge-8/0/8 unit 11 encapsulation vlan-ccc
set interfaces ge-8/0/8 unit 11 layer2-policer input-three-color abc
set interfaces ge-8/0/8 unit 11 family ccc
set interfaces ge-8/0/9 enable
set interfaces ge-8/0/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-8/0/9 unit 0 family mpls
set interfaces ae0 unit 0 family inet
set interfaces lo0 unit 0 family inet address 3.3.3.3/32
set interfaces ge-8/0/8 flexible-vlan-tagging
set interfaces ge-8/0/8 unit 11 vlan-id 2000
set interfaces ge-8/0/8 unit 11 input-vlan-map swap
set interfaces ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-8/0/8 unit 11 output-vlan-map swap
set routing-options router-id 3.3.3.3
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word
set protocols oam ethernet connectivity-fault-management maintenance-domain md level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 1 interface ge-8/0/8.11
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 1 direction up

```

```

set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 1 remote-mep 2
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k

```

### **Configuring Router PE1**

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

To configure Router PE1:

1. Configure the interfaces.

```

[edit]
user@PE1# edit interfaces
[edit interfaces]
user@PE1# set ge-5/0/4 encapsulation flexible-ethernet-services
user@PE1# set ge-5/0/4 unit 11 encapsulation vlan-ccc
user@PE1# set ge-5/0/4 unit 11 layer2-policer input-three-color abc
user@PE1# set ge-5/0/4 unit 11 family ccc
user@PE1# set ge-5/1/9 enable
user@PE1# set ge-5/1/9 unit 0 family inet address 12.1.1.1/24
user@PE1# set ge-5/1/9 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 4.4.4.4/32

```

2. Configure the VLAN.

```

[edit interfaces]
user@PE1# set ge-5/0/4 flexible-vlan-tagging
user@PE1# set ge-5/0/4 unit 11 vlan-id 2000
user@PE1# set ge-5/0/4 unit 11 input-vlan-map swap
user@PE1# set ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
user@PE1# set ge-5/0/4 unit 11 output-vlan-map swap

```

3. Configure the router identifier to identify the routing device.

```

[edit]
user@PE1# edit routing-options

```

```
[edit routing-options]
user@PE1# set router-id 4.4.4.4
```

#### 4. Configure MPLS, OSPF, and LDP protocols.

```
[edit]
user@PE1# edit protocols
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ldp interface all
user@PE1# set ldp interface fxp0.0 disable
```

#### 5. Configure the Layer 2 circuit.

```
[edit protocols]
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word
```

#### 6. Configure the MEP.

```
[edit protocols]
user@PE1# set oam ethernet connectivity-fault-management performance-monitoring
  delegate-server-processing
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md level 4
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma continuity-check interval 1s
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 interface ge-5/0/4.11
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 direction up
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 2 remote-mep 1
```

#### 7. Configure the firewall.

```
[edit]
```

```

user@PE1# edit firewall
[edit firewall]
user@PE1# set three-color-policer abc logical-interface-policer
user@PE1# set three-color-policer abc two-rate color-blind
user@PE1# set three-color-policer abc two-rate committed-information-rate 10m
user@PE1# set three-color-policer abc two-rate committed-burst-size 1500
user@PE1# set three-color-policer abc two-rate peak-information-rate 20m
user@PE1# set three-color-policer abc two-rate peak-burst-size 15k

```

8. Commit the configuration.

```

[edit]
user@PE1# commit

```

## Results

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

```

user@PE1# show interfaces
interfaces {
  ge-5/0/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 11 {
      encapsulation vlan-ccc;
      vlan-id 2000;
      input-vlan-map {
        swap;
        vlan-id 4094;
      }
      output-vlan-map swap;
      layer2-policer {
        input-three-color abc;
      }
      family ccc;
    }
  }
  ge-5/1/9 {
    enable;
    unit 0 {

```



```

        family inet {
            address 12.1.1.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 4.4.4.4/32;
        }
    }
}
}

```

user@PE1# **show protocols**

```

protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
    ldp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    l2circuit {
        neighbor 3.3.3.3 {
            interface ge-5/0/4.11 {
                virtual-circuit-id 1003;
                no-control-word;
            }
        }
    }
}

```



## Configuring Router PE2

### Step-by-Step Procedure

To configure Router PE2:

1. Configure the interfaces.

```
[edit]
user@PE2# edit interfaces
[edit interfaces]
user@PE2# set ge-8/0/8 encapsulation flexible-ethernet-services
user@PE2# set ge-8/0/8 unit 11 encapsulation vlan-ccc
user@PE2# set ge-8/0/8 unit 11 layer2-policer input-three-color abc
user@PE2# set ge-8/0/8 unit 11 family ccc
user@PE2# set ge-8/0/9 enable
user@PE2# set ge-8/0/9 unit 0 family inet address 12.1.1.1/24
user@PE2# set ge-8/0/9 unit 0 family mpls
user@PE2# set ae0 unit 0 family inet
user@PE2# set lo0 unit 0 family inet address 3.3.3.3/32
```

2. Configure the VLAN.

```
[edit interfaces]
user@PE2# set ge-8/0/8 flexible-vlan-tagging
user@PE2# set ge-8/0/8 unit 11 vlan-id 2000
user@PE2# set ge-8/0/8 unit 11 input-vlan-map swap
user@PE2# set ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
user@PE2# set ge-8/0/8 unit 11 output-vlan-map swap
```

3. Configure the router identifier to identify the routing device.

```
[edit]
user@PE2# edit routing-options
[edit routing-options]
user@PE2# set router-id 3.3.3.3
```

4. Configure MPLS, OSPF, and LDP protocols.

```
[edit]
user@PE2# edit protocols
[edit protocols]
```

```

user@PE2# set mpls interface all
user@PE2# set mpls interface fxp0.0 disable
user@PE2# set ospf area 0.0.0.0 interface all
user@PE2# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE2# set ldp interface all
user@PE2# set ldp interface fxp0.0 disable

```

5. Configure the Layer 2 circuit.

```

[edit protocols]
user@PE2# set l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id 1003
user@PE2# set l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word

```

6. Configure the MEP.

```

[edit protocols]
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md level 4
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
    maintenance-association ma continuity-check interval 1s
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
    maintenance-association ma mep 1 interface ge-8/0/8.11
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
    maintenance-association ma mep 1 direction up
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
    maintenance-association ma mep 1 remote-mep 2

```

7. Configure the firewall.

```

[edit]
user@PE2# edit firewall
[edit firewall]
user@PE2# set three-color-policer abc logical-interface-policer
user@PE2# set three-color-policer abc two-rate color-blind
user@PE2# set three-color-policer abc two-rate committed-information-rate 10m
user@PE2# set three-color-policer abc two-rate committed-burst-size 1500
user@PE2# set three-color-policer abc two-rate peak-information-rate 20m
user@PE2# set three-color-policer abc two-rate peak-burst-size 15k

```

## 8. Commit the configuration.

```
[edit]
user@PE2# commit
```

## Results

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

```
user@PE2# show interfaces
interfaces {
  ge-8/0/8 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 11 {
      encapsulation vlan-ccc;
      vlan-id 2000;
      input-vlan-map {
        swap;
        vlan-id 4094;
      }
      output-vlan-map swap;
      layer2-policer {
        input-three-color abc;
      }
      family ccc;
    }
  }
  ge-8/0/9 {
    unit 0 {
      family inet {
        address 12.1.1.2/24;
      }
      family mpls;
    }
  }
  ae0 {
    unit 0 {
      family inet;
    }
  }
  lo0 {
```

```

    unit 0 {
        family inet {
            address 3.3.3.3/32;
        }
    }
}

```

user@PE2# **show protocols**

```

protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  l2circuit {
    neighbor 4.4.4.4 {
      interface ge-8/0/8.11 {
        virtual-circuit-id 1003;
        no-control-word;
      }
    }
  }
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain md {
          level 4;
          maintenance-association ma {

```

```

        continuity-check {
            interval 1s;
        }
        mep 1 {
            interface ge-8/0/8.11;
            direction up;
            remote-mep 2;
        }
    }
}
}
}
}
}
}
}
}

```

```

user@PE2# show routing-options
routing-options {
    router-id 3.3.3.3;
}

```

```

user@PE2# show firewall
firewall {
    three-color-policer abc {
        logical-interface-policer;
        two-rate {
            color-blind;
            committed-information-rate 10m;
            committed-burst-size 1500;
            peak-information-rate 20m;
            peak-burst-size 15k;
        }
    }
}
}
}

```

### Verification

#### IN THIS SECTION

- [Viewing ETH-LM | 764](#)

To start monitoring the Ethernet frame loss, issue the **monitor ethernet loss-measurement maintenance-domain md maintenance-association ma mep 1** command. Frame loss is calculated by collecting the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs. The loss measurement statistics are retrieved as the output of the **monitor ethernet loss-measurement** command. You can also issue the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command to display ETH-LM statistics.

### Viewing ETH-LM

#### Purpose

View the ETH-LM statistics.

#### Action

From operational mode, enter the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command.

```
user@PE1> show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11
```

```
Interface name: ge-5/0/4.11 , Interface status: Active, Link status: Up
Maintenance domain name: md, Format: string, Level: 4
Maintenance association name: ma, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no
MEP identifier: 2, Direction: up, MAC address: 00:24:dc:9b:96:76
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : no
  Cross-connect CCM received                   : no
  RDI sent by some MEP                        : no
  Some remote MEP's MAC in error state         : no
Statistics:
  CCMs sent                                   : 36
  CCMs received out of sequence                : 0
  LBMs sent                                   : 0
  Valid in-order LBRs received                 : 0
  Valid out-of-order LBRs received             : 0
  LBRs received with corrupted data            : 0
  LBRs sent                                    : 0
  LTMs sent                                   : 0
  LTMs received                               : 0
  LTRs sent                                   : 0
  LTRs received                               : 0
```



```

Sequence number of next LTM request      : 0
1DMs sent                               : 0
Valid 1DMs received                      : 0
Invalid 1DMs received                    : 0
Out of sync 1DMs received                : 0
DMMs sent                               : 0
Valid DMMs received                     : 0
Invalid DMMs received                   : 0
DMRs sent                               : 0
Valid DMRs received                     : 0
Invalid DMRs received                   : 0
LMMs sent                               : 10
Valid LMMs received                     : 0
Invalid LMMs received                   : 0
LMRs sent                               : 0
Valid LMRs received                     : 10
Invalid LMRs received                   : 0
SLMs sent                               : 0
Valid SLMs received                     : 0
Invalid SLMs received                   : 0
SLRs sent                               : 0
Valid SLRs received                     : 0
Invalid SLRs received                   : 0
Remote MEP count: 1
Identifier    MAC address    State    Interface
    1        00:05:85:76:e5:30    ok    ge-5/0/4.11

```

### Meaning

The Ethernet interface details and statistics are displayed. This output indicates that the **ge-5/0/4.11** interface is active and its link status is **up**. Its maintenance domain name is **md** and its level is **4**. The MEP identifier of the **ge-5/0/4.11** interface is indicated as **2** and its direction is **up**. Under the statistics section, the output indicates that **10** LMMs were sent and **10** valid LMRs were received by the interface.

### SEE ALSO

[Configuring Ethernet Synthetic Loss Measurements | 796](#)

[Introduction to OAM Connectivity Fault Management \(CFM\) | 546](#)

## Example: Measuring Ethernet Frame Loss for Dual-Tagged LMM/LMR PDUs

### IN THIS SECTION

- Requirements | 766
- Overview and Topology | 766
- Configuration | 767
- Verification | 779

This example illustrates how to configure Ethernet frame loss measurement (ETH-LM) for dual-tagged Loss Measurement Message (LMM)/Loss Measurement Reply (LMR) protocol data units (PDUs). By configuring ETH-LM, you can measure the Ethernet frame loss that occur in your network.

### Requirements

This example uses the following hardware and software components:

- Two MX Series 5G Universal Routing Platforms with Rev-B Dense Port Concentrators (DPCs)
- Junos OS Release 14.2 or later

### Overview and Topology

Junos OS supports Ethernet frame loss measurement (ETH-LM) between maintenance association end points (MEPs) configured on Ethernet physical or logical interfaces on Rev-B Dense Port Concentrators (DPCs) in MX Series routers. Additionally, the Y.1731 functionality supports ETH-LM only for an end-to-end connection that uses Virtual Private Wire Service (VPWS). This example illustrates how to configure ETH-LM for dual tagged LMM/LMR PDUs with input and output VLAN map configured as **swap-swap**.

Figure 26 on page 766 shows the topology used in this example. VPWS service is configured between two MX Series routers, MX-PE1 and MX PE2.

Figure 26: VPWS Service Configured Between Two MX Series Routers



4 Level 4 UP MEP for Y1731 packets (MX Series client and MX Series server)

8042702

MX-PE1 router has two Ethernet interfaces, **ge-5/0/4** and **ge-5/1/9**. Virtual LAN (VLAN) is configured on **ge-5/0/4** and MPLS is configured on the **ge-5/1/9** interface. The **ge-5/0/4.11** interface is used to configure the Layer 2 virtual circuit with MX-PE2 router. The UP MEP, **mep 2**, is attached to the **ge-5/0/4.11** interface. The three-color policer firewall filter is also configured for the MX-PE1 router.

Similarly, MX-PE2 router has two Ethernet interfaces, **ge-8/0/8** and **ge-8/0/9**. Virtual LAN (VLAN) is configured on **ge-8/0/8** and MPLS is configured on the **ge-8/0/9** interface. The **ge-8/0/8.11** interface is used to configure the Layer 2 virtual circuit with MX-PE1 router. The UP MEP, **mep 1**, is attached to the **ge-8/0/8.11** interface. The three-color policer firewall filter is also configured for the MX-PE2 router.

### Configuration

#### IN THIS SECTION

- [Configuring Router PE1 | 769](#)
- [Configuring Router PE2 | 774](#)

### CLI Quick Configuration

To quickly configure ETH-LM for dual tagged LMM/LMR PDUs, copy the following commands, remove any line breaks, and then paste the commands into the CLI of each device.

On Router PE1:

```
[edit]
set interfaces ge-5/0/4 encapsulation flexible-ethernet-services
set interfaces ge-5/0/4 unit 11 encapsulation vlan-ccc
set interfaces ge-5/0/4 unit 11 layer2-policer input-three-color abc
set interfaces ge-5/0/4 unit 11 family ccc
set interfaces ge-5/1/9 enable
set interfaces ge-5/1/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-5/1/9 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 4.4.4.4/32
set interfaces ge-5/0/4 flexible-vlan-tagging
set interfaces ge-5/0/4 unit 11 vlan-tags outer 2000 inner 1000
set interfaces ge-5/0/4 unit 11 input-vlan-map swap-swap
set interfaces ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-5/0/4 unit 11 input-vlan-map inner-vlan-id 4093
set interfaces ge-5/0/4 unit 11 output-vlan-map swap-swap
set routing-options router-id 4.4.4.4
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```

set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word
set protocols oam ethernet connectivity-fault-management performance-monitoring delegate-server-processing
set protocols oam ethernet connectivity-fault-management maintenance-domain md level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 2 interface ge-5/0/4.11
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 2 direction up
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 2 remote-mep 1
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k

```

On Router PE2:

```

[edit]
set interfaces ge-8/0/8 encapsulation flexible-ethernet-services
set interfaces ge-8/0/8 unit 11 encapsulation vlan-ccc
set interfaces ge-8/0/8 unit 11 layer2-policer input-three-color abc
set interfaces ge-8/0/8 unit 11 family ccc
set interfaces ge-8/0/9 enable
set interfaces ge-8/0/9 unit 0 family inet address 12.1.1.1/24
set interfaces ge-8/0/9 unit 0 family mpls
set interfaces ae0 unit 0 family inet
set interfaces lo0 unit 0 family inet address 3.3.3.3/32
set interfaces ge-8/0/8 flexible-vlan-tagging
set interfaces ge-8/0/8 unit 11 vlan-tags outer 2000 inner 1000
set interfaces ge-8/0/8 unit 11 input-vlan-map swap-swap
set interfaces ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
set interfaces ge-8/0/8 unit 11 input-vlan-map inner-vlan-id 4093
set interfaces ge-8/0/8 unit 11 output-vlan-map swap-swap
set routing-options router-id 3.3.3.3
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

```

set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id 1003
set protocols l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word
set protocols oam ethernet connectivity-fault-management maintenance-domain md level 4
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma continuity-check interval 1s
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 1 interface ge-8/0/8.11
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 1 direction up
set protocols oam ethernet connectivity-fault-management maintenance-domain md maintenance-association
  ma mep 1 remote-mep 2
set firewall three-color-policer abc logical-interface-policer
set firewall three-color-policer abc two-rate color-blind
set firewall three-color-policer abc two-rate committed-information-rate 10m
set firewall three-color-policer abc two-rate committed-burst-size 1500
set firewall three-color-policer abc two-rate peak-information-rate 20m
set firewall three-color-policer abc two-rate peak-burst-size 15k

```

### Configuring Router PE1

#### Step-by-Step Procedure

To configure Router PE1:

1. Configure the interfaces.

```

[edit]
user@PE1# edit interfaces
[edit interfaces]
user@PE1# set ge-5/0/4 encapsulation flexible-ethernet-services
user@PE1# set ge-5/0/4 unit 11 encapsulation vlan-ccc
user@PE1# set ge-5/0/4 unit 11 layer2-policer input-three-color abc
user@PE1# set ge-5/0/4 unit 11 family ccc
user@PE1# set ge-5/1/9 enable
user@PE1# set ge-5/1/9 unit 0 family inet address 12.1.1.1/24
user@PE1# set ge-5/1/9 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 4.4.4.4/32

```

2. Configure the VLAN.

```

[edit interfaces]
user@PE1# set ge-5/0/4 flexible-vlan-tagging

```

```

user@PE1# set ge-5/0/4 unit 11 vlan-tags outer 2000 inner 1000
user@PE1# set ge-5/0/4 unit 11 input-vlan-map swap-swap
user@PE1# set ge-5/0/4 unit 11 input-vlan-map vlan-id 4094
user@PE1# set ge-5/0/4 unit 11 input-vlan-map inner-vlan-id 4093
user@PE1# set ge-5/0/4 unit 11 output-vlan-map swap-swap

```

3. Configure the router identifier to identify the routing device.

```

[edit]
user@PE1# edit routing-options
[edit routing-options]
user@PE1# set router-id 4.4.4.4

```

4. Configure MPLS, OSPF, and LDP protocols.

```

[edit]
user@PE1# edit protocols
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ldp interface all
user@PE1# set ldp interface fxp0.0 disable

```

5. Configure the Layer 2 circuit.

```

[edit protocols]
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 virtual-circuit-id 1003
user@PE1# set l2circuit neighbor 3.3.3.3 interface ge-5/0/4.11 no-control-word

```

6. Configure the MEP.

```

[edit protocols]
user@PE1# set oam ethernet connectivity-fault-management performance-monitoring
    delegate-server-processing
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md level 4
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
    maintenance-association ma continuity-check interval 1s

```

```

user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
maintenance-association ma mep 2 interface ge-5/0/4.11
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
maintenance-association ma mep 2 direction up
user@PE1# set oam ethernet connectivity-fault-management maintenance-domain md
maintenance-association ma mep 2 remote-mep 1

```

## 7. Configure the firewall.

```

[edit]
user@PE1# edit firewall
[edit firewall]
user@PE1# set three-color-policer abc logical-interface-policer
user@PE1# set three-color-policer abc two-rate color-blind
user@PE1# set three-color-policer abc two-rate committed-information-rate 10m
user@PE1# set three-color-policer abc two-rate committed-burst-size 1500
user@PE1# set three-color-policer abc two-rate peak-information-rate 20m
user@PE1# set three-color-policer abc two-rate peak-burst-size 15k

```

## 8. Commit the configuration.

```

[edit]
user@PE1# commit

```

## Results

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

```

user@PE1# show interfaces
interfaces {
  ge-5/0/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 11 {
      encapsulation vlan-ccc;
      vlan-tags outer 2000 inner 1000;
      input-vlan-map {
        swap-swap;
        vlan-id 4094;
      }
    }
  }
}

```

```

        inner-vlan-id 4093;
    }
    output-vlan-map swap-swap;
    layer2-policer {
        input-three-color abc;
    }
    family ccc;
}
}
ge-5/1/9 {
    enable;
    unit 0 {
        family inet {
            address 12.1.1.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 4.4.4.4/32;
        }
    }
}
}
}

```

user@PE1# **show protocols**

```

protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}
ldp {

```



```

interface all;
interface fxp0.0 {
    disable;
}
}
l2circuit {
    neighbor 3.3.3.3 {
        interface ge-5/0/4.11 {
            virtual-circuit-id 1003;
            no-control-word;
        }
    }
}
oam {
    ethernet {
        connectivity-fault-management {
            performance-monitoring {
                delegate-server-processing;
            }
            maintenance-domain md {
                level 4;
                maintenance-association ma {
                    continuity-check {
                        interval 1s;
                    }
                    mep 2 {
                        interface ge-5/0/4.11;
                        direction up;
                        remote-mep 1;
                    }
                }
            }
        }
    }
}
}
}

```

```

user@PE1# show routing-options
routing-options {
    router-id 4.4.4.4;
}

```

```

user@PE1# show firewall

```

```

firewall {
  three-color-policer abc {
    logical-interface-policer;
    two-rate {
      color-blind;
      committed-information-rate 10m;
      committed-burst-size 1500;
      peak-information-rate 20m;
      peak-burst-size 15k;
    }
  }
}

```

## Configuring Router PE2

### Step-by-Step Procedure

To configure Router PE2:

1. Configure the interfaces.

```

[edit]
user@PE2# edit interfaces
[edit interfaces]
user@PE2# set ge-8/0/8 encapsulation flexible-ethernet-services
user@PE2# set ge-8/0/8 unit 11 encapsulation vlan-ccc
user@PE2# set ge-8/0/8 unit 11 layer2-policer input-three-color abc
user@PE2# set ge-8/0/8 unit 11 family ccc
user@PE2# set ge-8/0/9 enable
user@PE2# set ge-8/0/9 unit 0 family inet address 12.1.1.1/24
user@PE2# set ge-8/0/9 unit 0 family mpls
user@PE2# set ae0 unit 0 family inet
user@PE2# set lo0 unit 0 family inet address 3.3.3.3/32

```

2. Configure the VLAN.

```

[edit interfaces]
user@PE2# set ge-8/0/8 flexible-vlan-tagging
user@PE2# set ge-8/0/8 unit 11 vlan-tags outer 2000 inner 1000
user@PE2# set ge-8/0/8 unit 11 input-vlan-map swap-swap
user@PE2# set ge-8/0/8 unit 11 input-vlan-map vlan-id 4094
user@PE2# set ge-8/0/8 unit 11 input-vlan-map inner-vlan-id 4093
user@PE2# set ge-8/0/8 unit 11 output-vlan-map swap-swap

```

3. Configure the router identifier to identify the routing device.

```
[edit]
user@PE2# edit routing-options
[edit routing-options]
user@PE2# set router-id 3.3.3.3
```

4. Configure MPLS, OSPF, and LDP protocols.

```
[edit]
user@PE2# edit protocols
[edit protocols]
user@PE2# set mpls interface all
user@PE2# set mpls interface fxp0.0 disable
user@PE2# set ospf area 0.0.0.0 interface all
user@PE2# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE2# set ldp interface all
user@PE2# set ldp interface fxp0.0 disable
```

5. Configure the Layer 2 circuit.

```
[edit protocols]
user@PE2# set l2circuit neighbor 4.4.4.4 interface ge-8/0/8.11 virtual-circuit-id 1003
user@PE2# set l2circuit neighbor 3.3.3.3 interface ge-8/0/8.11 no-control-word
```

6. Configure the MEP.

```
[edit protocols]
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md level 4
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma continuity-check interval 1s
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 1 interface ge-8/0/8.11
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 1 direction up
user@PE2# set oam ethernet connectivity-fault-management maintenance-domain md
  maintenance-association ma mep 1 remote-mep 2
```

## 7. Configure the firewall.

```
[edit]
user@PE2# edit firewall
[edit firewall]
user@PE2# set three-color-policer abc logical-interface-policer
user@PE2# set three-color-policer abc two-rate color-blind
user@PE2# set three-color-policer abc two-rate committed-information-rate 10m
user@PE2# set three-color-policer abc two-rate committed-burst-size 1500
user@PE2# set three-color-policer abc two-rate peak-information-rate 20m
user@PE2# set three-color-policer abc two-rate peak-burst-size 15k
```

## 8. Commit the configuration.

```
[edit]
user@PE2# commit
```

## Results

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

```
user@PE2# show interfaces
interfaces {
  ge-8/0/8 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 11 {
      encapsulation vlan-ccc;
      vlan-tags outer 2000 inner 1000;
      input-vlan-map {
        swap-swap;
        vlan-id 4094;
        inner-vlan-id 4093;
      }
      output-vlan-map swap-swap;
      layer2-policer {
        input-three-color abc;
      }
      family ccc;
    }
  }
}
```

```

ge-8/0/9 {
  unit 0 {
    family inet {
      address 12.1.1.2/24;
    }
    family mpls;
  }
}
ae0 {
  unit 0 {
    family inet;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 3.3.3.3/32;
    }
  }
}
}

```

user@PE2# **show protocols**

```

protocols {
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface all;
      interface fxp0.0 {
        disable;
      }
    }
  }
  ldp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}

```

```

l2circuit {
  neighbor 4.4.4.4 {
    interface ge-8/0/8.11 {
      virtual-circuit-id 1003;
      no-control-word;
    }
  }
}
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md {
        level 4;
        maintenance-association ma {
          continuity-check {
            interval 1s;
          }
          mep 1 {
            interface ge-8/0/8.11;
            direction up;
            remote-mep 2;
          }
        }
      }
    }
  }
}
}

```

**user@PE2# show routing-options**

```

routing-options {
  router-id 3.3.3.3;
}

```

**user@PE2# show firewall**

```

firewall {
  three-color-policer abc {
    logical-interface-policer;
    two-rate {
      color-blind;
      committed-information-rate 10m;
      committed-burst-size 1500;
      peak-information-rate 20m;
    }
  }
}

```

```

        peak-burst-size 15k;
    }
}
}

```

## Verification

### IN THIS SECTION

- [Viewing ETH-LM | 779](#)

To start the Ethernet frame loss measurement session, issue the **monitor ethernet loss-measurement maintenance-domain md maintenance-association ma mep 1** command. Frame loss is calculated by collecting the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs. The loss measurement statistics are retrieved as the output of the **monitor ethernet loss-measurement** command. You can also issue the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command to display ETH-LM statistics.

### Viewing ETH-LM

#### Purpose

View the ETH-LM statistics.

#### Action

From operational mode, enter the **show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11** command.

```
user@PE1> show oam ethernet connectivity-fault-management interfaces detail ge-5/0/4.11
```

```

Interface name: ge-5/0/4.11 , Interface status: Active, Link status: Up
Maintenance domain name: md, Format: string, Level: 4
Maintenance association name: ma, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no
MEP identifier: 2, Direction: up, MAC address: 00:24:dc:9b:96:76
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : no

```

```

Cross-connect CCM received           : no
RDI sent by some MEP                 : no
Some remote MEP's MAC in error state : no
Statistics:
CCMs sent                           : 59
CCMs received out of sequence        : 0
LBMs sent                           : 0
Valid in-order LBRs received         : 0
Valid out-of-order LBRs received     : 0
LBRs received with corrupted data    : 0
LBRs sent                           : 0
LTMs sent                           : 0
LTMs received                       : 0
LTRs sent                           : 0
LTRs received                       : 0
Sequence number of next LTM request  : 0
1DMs sent                           : 0
Valid 1DMs received                  : 0
Invalid 1DMs received                : 0
Out of sync 1DMs received            : 0
DMMs sent                           : 0
Valid DMMs received                  : 0
Invalid DMMs received                : 0
DMRs sent                           : 0
Valid DMRs received                  : 0
Invalid DMRs received                : 0
LMMs sent                           : 10
Valid LMMs received                  : 0
Invalid LMMs received                : 0
LMRs sent                           : 0
Valid LMRs received                  : 10
Invalid LMRs received                : 0
SLMs sent                           : 0
Valid SLMs received                  : 0
Invalid SLMs received                : 0
SLRs sent                           : 0
Valid SLRs received                  : 0
Invalid SLRs received                : 0
Remote MEP count: 1
Identifier  MAC address  State  Interface
    1      00:05:85:76:e5:30  ok    ge-5/0/4.11

```

## Meaning



The Ethernet interface details and statistics are displayed. This output indicates that the **ge-5/0/4.11** interface is active and its link status is **up**. Its maintenance domain name is **md** and its level is **4**. The MEP identifier of the **ge-5/0/4.11** interface is indicated as **2** and its direction is **up**. Under the statistics section, the output indicates that **10** LMMs were sent and **10** valid LMRs were received by the interface.

#### SEE ALSO

[Configuring Ethernet Synthetic Loss Measurements | 796](#)

[Introduction to OAM Connectivity Fault Management \(CFM\) | 546](#)

#### RELATED DOCUMENTATION

[ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)

[Configuring Ethernet Synthetic Loss Measurements | 796](#)

## Configuring an Iterator Profile

#### IN THIS SECTION

- [Configuring an Iterator Profile | 782](#)
- [Verifying the Configuration of an Iterator Profile | 785](#)
- [Managing Iterator Statistics | 788](#)
- [Configuring a Remote MEP with an Iterator Profile | 794](#)

Use this topic to configure an iterator profile that periodically transmits SLA measurement packets for delay and loss measurement. You can also view and clear the iterator statistics, and configure a remote MEP with an iterator profile.

## Configuring an Iterator Profile

You can create an iterator profile with its parameters to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for delay measurement or loss measurement.

To create an iterator profile:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management performance-monitoring
```

2. Configure the SLA measurement monitoring iterator:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# edit sla-iterator-profiles
```

3. Configure an iterator profile—for example, i1:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
user@host# set i1
```

4. (Optional) Configure the cycle time, which is the amount of time (in milliseconds) between back-to-back transmission of SLA frames for one connection, with values from 10 through 3,600,000. The default value is 1000 ms.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set cycle-time cycle-time-value
```

5. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value), with values from 1 through 2000. The default value is 2000.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set iteration-period iteration-period-value
```

6. Configure the measurement type as loss measurement, statistical frame-loss measurement, or two-way delay measurement.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set measurement-type (loss | statistical-frame-loss | two-way-delay)
```

7. (Optional) Configure the calculation weight for delay with values from 1 through 65,535. The default value is 1 (applicable only for two-way delay measurement).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set calculation-weight delay delay-value
```

8. (Optional) Configure the calculation weight for delay variation with values from 1 through 65,535. The default value is 1 (applicable only for two-way delay measurement).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set calculation-weight delay-variation delay-variation-value
```

9. (Optional) Configure the threshold value for average frame delay, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay is exceeded, an SNMP trap is generated for ETH-DM. The range is from 1 through 4294967295 microseconds.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set avg-fd-twoway-threshold avg-fd-twoway-threshold-value
```

10. (Optional) Configure the threshold value for average frame delay variation, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay variation is exceeded, an SNMP trap is generated for ETH-DM. The range is from 1 through 4294967295 microseconds.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set avg-ifdv-twoway-threshold avg-ifdv-twoway-threshold-value
```

11. (Optional) Configure the threshold value for average frame loss ratio, in milli-percent, in the upward or forward direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average forward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. The range is from 1 through 100000 milli-percent.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set avg-flr-forward-threshold avg-flr-forward-threshold-value
```

12. (Optional) Configure the threshold value for average frame loss ratio, in milli-percent, in the backward or downstream direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average backward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. The range is from 1 through 100000 milli-percent.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set avg-flr-backward-threshold avg-flr-backward-threshold-value
```

13. Configure the **disable** statement to stop the iterator (that is, disable the iterator profile).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set disable
```

14. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
user@host# show i1
  cycle-time cycle-time-value;
  iteration-period iteration-period-value;
  measurement-type (loss | two-way-delay);
  avg-fd-twoway-threshold avg-fd-twoway-threshold-value;
  avg-ifdv-twoway-threshold avg-ifdv-twoway-threshold-value;
  avg-flr-forward-threshold avg-flr-forward-threshold-value;
  avg-flr-backward-threshold avg-flr-backward-threshold-value;
  calculation-weight {
    delay delay-weight;
    delay-variation delay-variation-weight;
  }
  calculation-weight {
```

```

delay delay-weight;
delay-variation delay-variation-weight;
}

```

SEE ALSO

[Proactive Mode for SLA Measurement](#) | 700

## Verifying the Configuration of an Iterator Profile

### IN THIS SECTION

- [Displaying the Configuration of an Iterator Profile for Two-way Delay Measurement](#) | 785
- [Displaying the Configuration of an Iterator Profile for Loss Measurement](#) | 786
- [Displaying the Configuration of a Remote MEP with an Iterator Profile](#) | 787
- [Disabling an Iterator Profile](#) | 787

The following topics illustrate the configuration of an iterator profile for a two-way delay measurement, for loss measurement, and for a remote maintenance association end point (MEP). The topics also illustrate disabling an iterator profile with the **disable** statement for two-way measurement and deactivating an iterator profile with the **deactivate** command for a remote MEP.

### *Displaying the Configuration of an Iterator Profile for Two-way Delay Measurement*

#### Purpose

Display the configuration of an iterator profile for two-way delay measurement as configured in the [“Configuring an Iterator Profile” on page 782](#) topic with the following values:

- **profile-name**—i1
- **cycle-time**—1000 milliseconds
- **iteration-period**—2000 cycles per second
- **delay**—1
- **delay-variation**—1:

#### Action

To display information about the iterator profile, run the **show** command at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
user@host# show
i1 {
  cycle-time 1000;
  iteration-period 2000;
  measurement-type two-way-delay;
  calculation-weight {
    delay 1;
    delay-variation 1;
  }
}
```

### Meaning

The configuration for an iterator profile for two-way measurement is displayed as expected with set values.

### *Displaying the Configuration of an Iterator Profile for Loss Measurement*

#### Purpose

Display the configuration of an iterator profile for loss measurement as configured in the [“Configuring an Iterator Profile” on page 782](#) topic with the following values:

- **profile-name—12**
- **cycle-time—1000** milliseconds
- **iteration-period—2000** cycles per second

#### Action

To display information about the iterator profile, run the **show** command at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
user@host# show
12 {
  cycle-time 1000;
  iteration-period 2000;
  measurement-type loss;
}
```

### Meaning

The configuration for an iterator profile for loss measurement is displayed as expected with set values.

### *Displaying the Configuration of a Remote MEP with an Iterator Profile*

#### **Purpose**

Display the configuration of a remoteMEP as configured in the [“Configuring a Remote MEP with an Iterator Profile” on page 794](#) topic with the following values:

- **profile-name—i3**
- **maintenance-domain—default-1**
- **maintenance-association—1**
- **short-name-format—2octet**
- **mep—1**
- **remote-mep—1**
- **data-tlv-size—1**
- **iteration-count—1**
- **priority—1**

#### **Action**

To display information about the remote MEP, run the **show** command at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain default-1 maintenance association ma1 mep 1 remote-mep 1]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain default-1 maintenance
association 1 short-name-format 2octet mep 1 remote-mep 1]
user@host# show
sla-iterator-profile i3 {
  data-tlv-size 1;
  iteration-count 1;
  priority 1;
}
```

#### **Meaning**

The configuration for a remote MEP for two-way measurement is displayed as expected with set values.

### *Disabling an Iterator Profile*

#### **Purpose**

To disable an iterator profile for two-way delay measurement and for a remote MEP.

#### **Action**

- To disable an iterator profile (for example, i1) with the **disable** configuration command for two-way measurement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles i1]
user@host# disable
```

- To disable an iterator profile for a remote MEP (for example, i2) with the **deactivate** configuration command at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain default-1 maintenance association ma1 mep 1 remote-mep 1]** hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain default-1 maintenance
association ma1 mep 1 remote-mep 1]
user@host# deactivate sla-iterator-profile i2
```

SEE ALSO

[Proactive Mode for SLA Measurement | 700](#)

## Managing Iterator Statistics

### IN THIS SECTION

- [Displaying Iterator Statistics | 788](#)
- [Clearing Iterator Statistics | 794](#)

### *Displaying Iterator Statistics*

#### Purpose

Retrieve and display iterator statistics.

Multiple iterators can be associated with a remote MEP. However, by default, only one result pertaining to one iterator profile is displayed.

#### Action

- To display the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1** (here, the iterator profile **i1** is configured for two-way delay measurement):



```
user@host> show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator
i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1
```

```
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i1, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent : 23898520
DMM skipped for threshold hit : 11000
DMM skipped for threshold hit window : 0
DMR received : 23851165
DMR out of sequence : 1142
DMR received with invalid time stamps : 36540
Average two-way delay : 129 usec
Average two-way delay variation : 15 usec
Average one-way forward delay variation : 22 usec
Average one-way backward delay variation : 22 usec
Weighted average two-way delay : 134 usec
Weighted average two-way delay variation : 8 usec
Weighted average one-way forward delay variation : 6 usec
Weighted average one-way backward delay variation : 2 usec
```

Output fields are listed in the approximate order in which they appear.

**Table 95: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields**

| Output Field Name              | Output Field Description              |
|--------------------------------|---------------------------------------|
| <b>Maintenance domain</b>      | Maintenance domain name.              |
| <b>Level</b>                   | Maintenance domain level configured.  |
| <b>Maintenance association</b> | Maintenance association name.         |
| <b>Local MEP id</b>            | Numeric identifier of the local MEP.  |
| <b>Remote MEP id</b>           | Numeric identifier of the remote MEP. |

Table 95: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields (*continued*)

| Output Field Name                            | Output Field Description                                                                                     |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Remote MAC address</b>                    | Unicast MAC address of the remote MEP.                                                                       |
| <b>Iterator name</b>                         | Name of iterator.                                                                                            |
| <b>Iterator Id</b>                           | Numeric identifier of the iterator.                                                                          |
| <b>Iterator cycle time</b>                   | Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection |
| <b>Iteration period</b>                      | Maximum number of cycles per iteration                                                                       |
| <b>Iterator status</b>                       | Current status of iterator whether running or stopped.                                                       |
| <b>Infinite iterations</b>                   | Status of iteration as infinite or finite.                                                                   |
| <b>Counter reset time</b>                    | Date and time when the counter was reset.                                                                    |
| <b>Reset reason</b>                          | Reason to reset counter.                                                                                     |
| <b>Delay weight</b>                          | Calculation weight of delay.                                                                                 |
| <b>Delay variation weight</b>                | Calculation weight of delay variation.                                                                       |
| <b>DMM sent</b>                              | Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.                             |
| <b>DMM skipped for threshold hit</b>         | Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.                      |
| <b>DMM skipped for threshold hit window</b>  | Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.      |
| <b>DMR received</b>                          | Number of delay measurement reply (DMR) frames received.                                                     |
| <b>DMR out of sequence</b>                   | Total number of DMR out of sequence packets received.                                                        |
| <b>DMR received with invalid time stamps</b> | Total number of DMR frames received with invalid timestamps.                                                 |
| <b>Average two-way delay</b>                 | Average two-way frame delay for the statistics displayed.                                                    |

Table 95: Displaying Iterator Statistics for Ethernet Delay Measurement Output Fields (*continued*)

| Output Field Name                                        | Output Field Description                                                                        |
|----------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| <b>Average two-way delay variation</b>                   | Average two-way “frame jitter” for the statistics displayed.                                    |
| <b>Average one-way forward delay variation</b>           | Average one-way forward delay variation for the statistics displayed in microseconds.           |
| <b>Average one-way backward delay variation</b>          | Average one-way backward delay variation for the statistics displayed in microseconds.          |
| <b>Weighted average two-way delay</b>                    | Weighted average two-way delay for the statistics displayed in microseconds.                    |
| <b>Weighted average two-way delay variation</b>          | Weighted average two-way delay variation for the statistics displayed in microseconds.          |
| <b>Weighted average one-way forward delay variation</b>  | Weighted average one-way forward delay variation for the statistics displayed in microseconds.  |
| <b>Weighted average one-way backward delay variation</b> | Weighted average one-way backward delay variation for the statistics displayed in microseconds. |

- To display the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1** (here, the iterator profile **i1** is configured for loss measurement):

```
user@host> show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator
i2 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1
```

```
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i2, Iterator Id: 2
Iterator cycle time: 1000ms, Iteration period: 2000 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:25 ago)
Reset reason: Adjacency flap

Iterator loss measurement statistics:
```

```

LMM sent : 238970
LMM skipped for threshold hit : 60
LMM skipped for threshold hit window : 0
LMR received : 238766
LMR out of sequence : 43

Accumulated transmit statistics:
Near-end (CIR) : 0
Far-end (CIR) : 0
Near-end (EIR) : 0
Far-end (EIR) : 0

Accumulated loss statistics:
Near-end (CIR) : 0 (0.00%)
Far-end (CIR) : 0 (0.00%)
Near-end (EIR) : 0 (0.00%)
Far-end (EIR) : 0 (0.00%)

Last loss measurement statistics:
Near-end (CIR) : 0
Far-end (CIR) : 0
Near-end (EIR) : 0
Far-end (EIR) : 0

```

Output fields are listed in the approximate order in which they appear.

**Table 96: Displaying Iterator Statistics for Ethernet Loss Measurement Output Fields**

| Output Field Name              | Output Field Description               |
|--------------------------------|----------------------------------------|
| <b>Maintenance domain</b>      | Maintenance domain name.               |
| <b>Level</b>                   | Maintenance domain level configured.   |
| <b>Maintenance association</b> | Maintenance association name.          |
| <b>Local MEP id</b>            | Numeric identifier of the local MEP.   |
| <b>RemoteMEP identifier</b>    | Numeric identifier of the remote MEP.  |
| <b>Remote MAC address</b>      | Unicast MAC address of the remote MEP. |
| <b>Iterator name</b>           | Name of iterator.                      |
| <b>Iterator Id</b>             | Numeric identifier of the iterator.    |

Table 96: Displaying Iterator Statistics for Ethernet Loss Measurement Output Fields (*continued*)

| Output Field Name                           | Output Field Description                                                                                     |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Iterator cycle time</b>                  | Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection |
| <b>Iteration period</b>                     | Maximum number of cycles per iteration                                                                       |
| <b>Iterator status</b>                      | Current status of iterator whether running or stopped.                                                       |
| <b>Infinite iterations</b>                  | Status of iteration as infinite or finite.                                                                   |
| <b>Counter reset time</b>                   | Date and time when the counter was reset.                                                                    |
| <b>Reset reason</b>                         | Reason to reset counter.                                                                                     |
| <b>LMM sent</b>                             | Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.                    |
| <b>LMM skipped for threshold hit</b>        | Number of LMM frames sent to the peer MEP in this session skipped during threshold hit.                      |
| <b>LMM skipped for threshold hit window</b> | Number of LMM frames sent to the peer MEP in this session skipped during the last threshold hit window.      |
| <b>LMR received</b>                         | Number of LMRs frames received.                                                                              |
| <b>LMR out of sequence</b>                  | Total number of LMR out of sequence packets received.                                                        |
| <b>Near-end (CIR)</b>                       | Frame loss associated with ingress data frames for the statistics displayed.                                 |
| <b>Far-end (CIR)</b>                        | Frame loss associated with egress data frames for the statistics displayed.                                  |
| <b>Near-end (EIR)</b>                       | Frame loss associated with ingress data frames for the statistics displayed.                                 |
| <b>Far-end (EIR)</b>                        | Frame loss associated with egress data frames for the statistics displayed.                                  |

SEE ALSO

[Proactive Mode for SLA Measurement | 700](#)
[show oam ethernet connectivity-fault-management sla-iterator-statistics | 1938](#)

```
| clear oam ethernet connectivity-fault-management sla-iterator-statistics
```

### Clearing Iterator Statistics

#### Purpose

Clear iterator statistics.

Multiple iterators can be associated with remote MEP. However, by default, only one result pertaining to one iterator profile can be cleared.

#### Action

- To clear the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1**:

```
user@host> clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i1
maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1
```

- To clear the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** and within the maintenance domain **default-1**:

```
user@host> clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i2
maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1
```

#### SEE ALSO

[Proactive Mode for SLA Measurement | 700](#)

[show oam ethernet connectivity-fault-management sla-iterator-statistics | 1938](#)

[clear oam ethernet connectivity-fault-management sla-iterator-statistics](#)

#### SEE ALSO

[Proactive Mode for SLA Measurement | 700](#)

### Configuring a Remote MEP with an Iterator Profile

You can associate a remote maintenance association end point (MEP) with more than one iterator profile.

To configure a remote MEP with an iterator profile:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id
```

2. Configure the remote MEP with values from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id]
user@host# set remote-mep remote-mep-id
```

3. Set the iterator profile.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep remote-mep-id]
user@host# set sla-iterator-profile profile-name
```

4. (Optional) Set the size of the data TLV portion of the Y.1731 data frame with values from 1 through 1400 bytes. The default value is 1.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
user@host# set data-tlv-size size
```

5. (Optional) Set the iteration count, which indicates the number of iterations for which this connection should partake in the iterator for acquiring SLA measurements, with values from 1 through 65,535. The default value is 0 (that is, infinite iterations).

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
user@host# set iteration-count count-value
```

6. (Optional) Set the priority, which is the **vlan-pcp** value that is sent in the Y.1731 data frames, with values from 0 through 7. The default value is 0.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
user@host# set priority priority-value
```

7. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id remote-mep remote-mep-id]
user@host# show
sla-iterator-profile profile-name {
  data-tlv-size size;
  iteration-count count-value;
  priority priority-value;
}
```

#### SEE ALSO

[Proactive Mode for SLA Measurement | 700](#)

[remote-mep | 1074](#)

#### RELATED DOCUMENTATION

[ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)

[Configuring Ethernet Frame Delay Measurement Sessions | 709](#)

## Configuring Ethernet Synthetic Loss Measurements

#### IN THIS SECTION

- [Guidelines for Configuring ETH-SLM | 797](#)
- [Starting a Proactive ETH-SLM Session | 798](#)
- [Starting an On-Demand ETH-SLM Session | 803](#)
- [Managing ETH-SLM Statistics and ETH-SLM Frame Counts | 804](#)
- [Troubleshooting Failures with ETH-SLM | 809](#)

Use this topic to understand the guidelines for configuring synthetic loss measurement and how to start a synthetic loss measurement session. There are two types of synthetic loss measurement sessions: proactive and On-Demand. This topic describes both. Also, the topic shows you how to view and clear the synthetic loss measurement statistics and how to troubleshoot failures with SLM.



## Guidelines for Configuring ETH-SLM

Keep the following points in mind when you configure the ETH-SLM functionality:

- The monitoring application for Ethernet OAM is initiated in the master Routing Engine. When a stateful switchover process occurs, the monitoring application is disabled. For on-demand ETH-SLM, graceful Routing Engine switchover (GRES) support is not applicable. For proactive ETH-SLM, the service-level agreement (SLA) iterators are restored during a stateful switchover process. If the adjacencies do not time out, the ETH-SLM statistics are preserved and proactive ETH-SLM supports GRES.
- ETH-SLM is initiated only when the MEP session is up. Unified in-service software upgrade (ISSU) support for ETH-SLM depends on the unified ISSU support for CFM. For CFM, unified ISSU is supported using the loss threshold TLV to avoid CFM connectivity loss during the upgrade. The receiving or the destination MEP increases the threshold time during the termination of sessions. If you start a unified ISSU operation when on-demand ETH-SLM is in progress, the SLM request and reply messages are lost at the local Packet Forwarding Engine.

When an on-demand ETH-SLM is requested, if the local source MEP undergoes a unified ISSU, a message is displayed stating that the MEP is undergoing a unified ISSU. If the remote MEP is undergoing a unified ISSU (detected through the loss threshold TLV), a message is displayed stating that the remote MEP is undergoing a unified ISSU. Also, if it is not possible to identify whether unified ISSU is in progress on a remote MEP, the SLM packets are lost at the system where unified ISSU is in progress and the loss calculation results do not provide a valid cause for the loss. Unified ISSU is not supported for both on-demand and proactive ETH-SLM.

- The maximum number of SLA iterator profiles that can be configured in the system is 255.
- ETH-SLM is not supported for virtual private LAN service (VPLS) (point-to-multipoint measurements are not supported). The ETH-SLM frames are not generated with multicast class 1 destination address. Similarly, ETH-SLM does not respond to ETH-SLM requests with multicast DA. ETH-SLM for VPLS for point-to-point Ethernet connection is supported using directed unicast destination MAC addresses, although point-to-multipoint topologies are not supported.
- A unicast destination address may be used in provisioned environments for point-to-point connections. However, it requires that the unicast destination address of the downstream MEP must have been configured on the MEP transmitting an alarm indication signal (AIS).
- ETH-SLM is not supported on downstream MEPs on label-switched interfaces (LSIs).
- ETH-SLM is supported on aggregated Ethernet (ae) interfaces
- The number of ETH-SLM sessions for proactive ETH-SLM that can be supported is limited to the total number of iterators that can be supported in the system. This limitation includes the iterator support for other measurement types such as loss, statistical frame loss, and two-way delay. A new iterator type, SLM, is added to support ETH-SLM. The total number of SLA iterators that you can configure in the system is equal to the total number of iterations supported in the system.
- For on-demand SLM, the minimum period between two SLM requests is 100 milliseconds.

- For proactive SLM, the minimum period between two SLM requests is 10 milliseconds for distributed mode and 100 milliseconds for non-distributed mode.
- ETH-SLM frames are always marked as drop-ineligible in compliance with the ITU-T Y.1731 standard.

#### SEE ALSO

[Ethernet Synthetic Loss Measurement Overview | 702](#)

[monitor ethernet synthetic-loss-measurement | 1463](#)

## Starting a Proactive ETH-SLM Session

### IN THIS SECTION

- [Configuring MEP Interfaces | 798](#)
- [Configuring an Iterator Profile for ETH-SLM | 799](#)
- [Associating the Iterator Profile with MEPs for ETH-SLM | 801](#)

To start a proactive Ethernet synthetic loss measurement (ETH-SLM) session, you must configure the Ethernet interfaces on maintenance association end points (MEPs) on which packets transmitted with synthetic frame loss need to be analyzed. You must then create an iterator profile to transmit service-level agreement (SLA) measurement packets for ETH-SLM and associate the local and remote MEPs with the profile.

### **Configuring MEP Interfaces**

Before you can start an Ethernet synthetic frame loss measurement session across an Ethernet service, you must configure two ACX Series routers to support ETH-SLM.

To configure an Ethernet interface on an ACX Series router to support ETH-SLM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

```
[edit interfaces]
interface {
  ethernet-interface-name {
    vlan-tagging;
    unit logical-unit-number {
```

```

        vlan-id vlan-id; # Both interfaces on this VLAN
    }
}

```

Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

```

[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md-name { # On both routers
        level number;
        maintenance-association ma-name { # On both routers
          continuity-check {
            interval 100ms;
            hold-interval 1;
          }
          mep mep-id { # Attach to VLAN interface
            auto-discovery;
            direction (up | down);
            interface interface-name;
            priority number;
          }
        }
      }
    }
  }
}

```

### Configuring an Iterator Profile for ETH-SLM

You can create an iterator profile with its parameters to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for synthetic loss measurement.

**NOTE:** ACX5048 and ACX5096 routers supports iterator cycle time of only 1 second and above.

To create an iterator profile:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management performance-monitoring
```

2. Configure the SLA measurement monitoring iterator:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# edit sla-iterator-profiles
```

3. Configure an iterator profile—for example, i1:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
user@host# set i1
```

4. (Optional) Configure the cycle time, which is the amount of time (in milliseconds) between back-to-back transmission of SLA frames for one connection, with a value from 10 through 3,600,000. The default value is 1000 ms.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set cycle-time cycle-time-value
```

5. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value), with a value from 1 through 2000. The default value is 2000.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set iteration-period iteration-period-value
```

6. Configure the measurement type as synthetic loss measurement.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set measurement-type slm
```

7. Configure the **disable** statement to stop the iterator (that is, disable the iterator profile).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
i1]
user@host# set disable
```

8. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
user@host# show i1
    cycle-time cycle-time-value;
    iteration-period iteration-period-value;
    measurement-type slm;
```

### ***Associating the Iterator Profile with MEPs for ETH-SLM***

You can associate a remote maintenance association end point (MEP) with more than one iterator profile.

To configure a remote MEP with an iterator profile:

1. In configuration mode, go to the following hierarchy level:

```
user@host# edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id
```

2. Configure the remote MEP ID with a value from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id]
user@host# set remote-mep remote-mep-id
```

3. Set the iterator profile.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep remote-mep-id]
user@host# set sla-iterator-profile profile-name
```

4. (Optional) Set the size of the data TLV portion of the Y.1731 data frame with a value from 1 through 1400 bytes. The default value is 1.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
```

```
user@host# set data-tlv-size size
```

5. (Optional) Set the iteration count, which indicates the number of iterations for which this connection should partake in the iterator for acquiring SLA measurements, with a value from 1 through 65,535. The default value is 0 (that is, infinite iterations).

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
 maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
user@host# set iteration-count count-value
```

6. (Optional) Set the priority, which is the **vlan-pcp** value that is sent in the Y.1731 data frames, with a value from 0 through 7. The default value is 0.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
 maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
user@host# set priority priority-value
```

7. Verify the configuration.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
 maintenance-association ma-name mep mep-id remote-mep remote-mep-id]
user@host# show
sla-iterator-profile profile-name {
  data-tlv-size size;
  iteration-count count-value;
  priority priority-value;
}
```

## SEE ALSO

[Ethernet Synthetic Loss Measurement Overview | 702](#)

[monitor ethernet synthetic-loss-measurement | 1463](#)

### Starting an On-Demand ETH-SLM Session

To start an on-demand Ethernet synthetic loss measurement (ETH-SLM) session, type the **monitor ethernet synthetic-loss-measurement one-way** command in operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```

```
ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0
  Synthetic Loss measurement statistics:
    SLM packets sent                : 100
    SLR packets received             : 100
    Accumulated SLM statistics:
      Local TXFC1 value              : 100
      Local RXFC1 value              : 100
      Last Received SLR frame TXFCf(tc) : 100
      Last Received SLR frame TXFCb(tc) : 100
      SLM Frame Loss:
        Frame Loss (far-end)          : 0 (0.00 %)
        Frame Loss (near-end)         : 0 (0.00 %)
```

**NOTE:** If you attempt to monitor delays to a nonexistent MAC address, you must press **Ctrl + C** to explicitly quit the **monitor ethernet synthetic-loss-measurement** command and return to the CLI command prompt.

SEE ALSO

- Ethernet Synthetic Loss Measurement Overview | 702
- monitor ethernet synthetic-loss-measurement | 1463

## Managing ETH-SLM Statistics and ETH-SLM Frame Counts

### IN THIS SECTION

- [Displaying ETH-SLM Statistics Only | 804](#)
- [Displaying ETH-SLM Statistics and Frame Counts | 805](#)
- [Displaying ETH-SLM Frame Counts for MEPs by Enclosing CFM Entity | 806](#)
- [Displaying ETH-SLM Frame Counts for MEPs by Interface or Domain Level | 807](#)
- [Clearing ETH-SLM Statistics and Frame Counts | 808](#)
- [Clearing Iterator Statistics | 808](#)

### *Displaying ETH-SLM Statistics Only*

#### Purpose

Display on-demand ETH-SLM statistics.

By default, the **show oam ethernet connectivity-fault-management synthetic-loss-statistics** command displays on-demand ETH-SLM statistics for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

#### Action

- To display the on-demand ETH-SLM statistics collected for MEPs belonging to maintenance association **ma1** within maintenance domain **md1**:

```
user@host> show oam ethernet connectivity-fault-management synthetic-loss-statistics  
maintenance-domain md1 maintenance-association ma1
```

- To display the on-demand ETH-SLM statistics collected for ETH-SLM sessions for the local MEP **201** belonging to maintenance association **ma2** within maintenance domain **md2**:

```
user@host> show oam ethernet connectivity-fault-management synthetic-loss-statistics  
maintenance-domain md2 maintenance-association ma2 local-mep 201
```

- To display the on-demand ETH-SLM statistics collected for ETH-SLM sessions from local MEPs belonging to maintenance association **ma3** within maintenance domain **md3** to the remote MEP **302**:

```
user@host> show oam ethernet connectivity-fault-management synthetic-loss-statistics  
maintenance-domain md3 maintenance-association ma3 remote-mep 302
```



## Meaning

The output displays on-demand ETH-SLM statistics for MEPs in the specified maintenance association within the specified maintenance domain. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management synthetic-loss-statistics**.

## SEE ALSO

| **show oam ethernet connectivity-fault-management synthetic-loss-statistics**

## Displaying ETH-SLM Statistics and Frame Counts

### Purpose

Display on-demand ETH-SLM statistics and ETH-SLM frame counts.

By default, the **show oam ethernet connectivity-fault-management mep-statistics** command displays on-demand ETH-SLM statistics and frame counts for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

### Action

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for MEPs in maintenance association **ma1** within maintenance domain **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain
md1 maintenance-association ma1
```

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for the local MEP **201** in maintenance association **ma2** within maintenance domain **md2**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain
md2 maintenance-association ma2 local-mep 201
```

- To display the on-demand ETH-SLM statistics and ETH-SLM frame counts for the local MEP in maintenance association **ma3** within maintenance domain **md3** that participates in an ETH-SLM session with the remote MEP **302**:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain
ma3 maintenance-association ma3 remote-mep 302
```

## Meaning

The output displays on-demand ETH-SLM statistics and ETH-SLM frame counts for MEPs in the specified maintenance association within the specified maintenance domain. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management mep-statistics**.

SEE ALSO

**show oam ethernet connectivity-fault-management mep-statistics**

### *Displaying ETH-SLM Frame Counts for MEPs by Enclosing CFM Entity*

#### **Purpose**

Display on-demand ETH-SLM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management mep-database** command displays CFM database information for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.

**NOTE:** At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display the ETH-SLM frame counts and not the MEP database details.

#### **Action**

- To display CFM database information (including ETH-SLM frame counts) for all MEPs in MA **ma1** within maintenance domain **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
ma1 maintenance-association ma1
```

- To display CFM database information (including ETH-SLM frame counts) only for the local MEP **201** in MA **ma1** within maintenance domain **md1**:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md2 maintenance-association ma2 local-mep 201
```

- To display CFM database information (including ETH-SLM frame counts) only for the remote MEP **302** in MA **ma3** within maintenance domain **md3**:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
ma3 maintenance-association ma3 remote-mep 302
```

#### **Meaning**

The output displays ETH-SLM frame counts for MEPs within a particular maintenance domain, or for a specific local or remote MEP. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management mep-database**.

SEE ALSO

| **show oam ethernet connectivity-fault-management mep-database**

### *Displaying ETH-SLM Frame Counts for MEPs by Interface or Domain Level*

#### **Purpose**

Display on-demand ETH-SLM frame counts for CFM maintenance association end points (MEPs).

By default, the **show oam ethernet connectivity-fault-management interfaces** command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-SLM, any frame counts are also displayed when you specify the **detail** or **extensive** command option.

**NOTE:** At the router attached to the initiator MEP, you can only display the ETH-SLM frame counts and not the MEP database details.

#### **Action**

- To display CFM database information (including ETH-SLM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
```

- To display CFM database information (including ETH-SLM frame counts) only for the MEPs attached to CFM-enabled router interface **ge-5/2/9.0**:

```
user@host> show oam ethernet connectivity-fault-management interfaces ge-5/2/9.0 detail
```

- To display CFM database information (including ETH-SLM frame counts) only for MEPs enclosed within CFM maintenance domains at level **6**:

```
user@host> show oam ethernet connectivity-fault-management interfaces level 6 detail
```

#### **Meaning**

The output displays ETH-SLM frame counts for MEPs for the specified interface. For details about the output of this command and the descriptions of the output fields, see **show oam ethernet connectivity-fault-management interfaces**.

SEE ALSO

| **show oam ethernet connectivity-fault-management interfaces**

### *Clearing ETH-SLM Statistics and Frame Counts*

#### **Purpose**

Clear the on-demand ETH-SLM statistics and ETH-SLM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

#### **Action**

- To clear the on-demand ETH-SLM statistics and ETH-SLM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

```
user@host> clear oam ethernet connectivity-fault-management synthetic-loss-measurement
```

- To clear the on-demand ETH-SLM statistics and ETH-SLM frame counts only for MEPs attached to the logical interface **ge-0/5.9.0**:

```
user@host> clear oam ethernet connectivity-fault-management synthetic-loss-measurement
ge-0/5/9.0
```

### *Clearing Iterator Statistics*

#### **Purpose**

Clear the existing iterator statistics and proactive ETH-SLM counters.

Multiple iterators can be associated with remote MEP. However, by default, only one result pertaining to one iterator profile can be cleared.

#### **Action**

- To clear the iterator statistics for remote MEP **1** and iterator profile **i1** with MEPs belonging to the maintenance association **ma1** within the maintenance domain **default-1**:

```
user@host> clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i1
maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1
```

- To clear the iterator statistics for remote MEP **1** and iterator profile **i2** with MEPs belonging to the maintenance association **ma1** within the maintenance domain **default-1**:

```
user@host> clear oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator i2
maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 1
```

SEE ALSO

[clear oam ethernet connectivity-fault-management synthetic-loss-measurement](#) | 1446

[show oam ethernet connectivity-fault-management synthetic-loss-statistics](#) | 1951

[show oam ethernet connectivity-fault-management interfaces](#) | 1879 (detail | extensive)

[show oam ethernet connectivity-fault-management mep-statistics](#) | 1916

[show oam ethernet connectivity-fault-management mep-database](#) | 1901

## Troubleshooting Failures with ETH-SLM

### Problem

**Description:** The Ethernet synthetic loss measurement (ETH-SLM) application is not working properly for calculation of frame loss using synthetic frames instead of data traffic

### Solution

Perform the following steps to analyze and debug any problems with the ETH-SLM functionality.

1. Ensure that ETH-SLM is configured (either proactive or on-demand) to initiate SLM frames. Verify the configuration settings.
2. Examine any failures that might have occurred in the CFM session for which the ETH-SLM feature is enabled. The CFM session must be in the up state for the ETH-SLM functionality to work correctly. Use the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *md-name* maintenance-association *ma-name* local-mep *mep-id* remote-mep *remote-mep-id*** command to verify whether the CFM session is in the up state.
3. If the MEP sessions are active, use the appropriate show command to verify the ETH-SLM statistics and to analyze if ETH-SLM frames are transmitted or received.
4. If the transmission of ETH-SLM frames does not happen correctly after you attempt all of the preceding troubleshooting steps, enable the tracing operations for Ethernet CFM by including the **traceoptions** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit protocols oam ethernet connectivity-fault-management]
traceoptions {
  file <filename> <files number> <match regular-expression microsecond-stamp>> <size size> <world-readable |
    no-world-readable>;
  flag <flag>;
  no-remote-trace;
}
```

SEE ALSO

---

[Ethernet Synthetic Loss Measurement Overview | 702](#)[monitor ethernet synthetic-loss-measurement | 1463](#)

---

## RELATED DOCUMENTATION

---

[ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)[Configuring Ethernet Frame Loss Measurement | 747](#)

---

## Ethernet Alarm Indication

### IN THIS SECTION

- [Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)
- [Ethernet Alarm Indication Signal Overview | 815](#)
- [Configuring ETH-AIS on a CFM MEP | 816](#)
- [Configuring Alarm Indication Signal on ACX Series Routers | 821](#)

Use the following to understand more about Ethernet alarm indication signal (ETH-AIS) and how to configure ETH-AIS on devices.

### Ethernet Alarm Indication Signal (ETH-AIS) Function Overview

#### IN THIS SECTION

- [Understanding ETH-AIS in a Maintenance Domain | 811](#)
- [Fault Detection in a Maintenance Domain | 811](#)
- [Terms Defined | 813](#)

Ethernet alarm indication signal (ETH-AIS) function enables a service provider deploying an Ethernet service to determine whether a connectivity fault exists at the provider's domain level or at a level below. When the fault occurs at the provider's domain level, the service provider addresses the fault, and when the fault occurs at a level below, the provider can either ignore the fault or contact the relevant authorities to address the fault.

The following sections explain ETH-AIS, few use cases which determine when to generate and propagate ETH-AIS packets, and associated terms in detail:

### ***Understanding ETH-AIS in a Maintenance Domain***

ITU-T developed Y.1731 as a recommendation for Operation, Administration, and Maintenance (OAM) functions and mechanisms for Ethernet-based networks, including OAM functions such as ETH-AIS, Ethernet locked signal (ETH-LCK), Ethernet test signal (ETH-Test), Ethernet automatic protection switching (ETH-APS), Ethernet maintenance communication channel (ETH-MCC), Ethernet experimental OAM (ETH-EXP), Ethernet vendor-specific OAM (ETH-VSP), and performance monitoring. For information about maintenance domain and related terms, see [“Terms Defined” on page 813](#).

According to the Y.1731 standards, a server MEP is a combined function of the server layer termination function and the server Ethernet services layer adaptation function. The server MEP notifies the Ethernet services (ETH) layer MEPs when it detects a failure. The server layer termination function then runs the OAM mechanisms specific to the server layer and the alarms are suppressed at the server layer by ETH-AIS.

Note that ETH-AIS is not applicable to Spanning Tree Protocol (STP) networks.

ETH-AIS enables you to suppress alarms when a fault condition is detected. Using ETH-AIS, a service provider can differentiate between faults at different levels.

ETH-AIS provides many advantages that include:

- Service providers need not raise alarms if there are lower-level failures.
- Service providers can provide a refund to their subscribers or avail a refund from their Internet provider based on service unavailability.

MX Series routers support ITU-T Y.1731 ETH-AIS to provide fault management for service providers who provide carrier Ethernet services using IEEE 802.1ag standard.


**NOTE:** MX Series Virtual Chassis does not support Ethernet alarm indication signal (ETH-AIS).

### ***Fault Detection in a Maintenance Domain***

In the scenario depicted in Figure 1 on page xyz, you have a service provider level and a customer level. Two service providers—*Operator-1* and *Operator-2*—are considered for illustration purposes. Assume that a fault occurs in *Operator-1* maintenance domain-level that has MEP-A and MEP-B at its maintenance domain-level boundaries. To notify the faults to a network management system and to avoid notification

of alarms from the customer level for the same fault, MEP-A and MEP-B transmit an alarm indication signal (AIS) on opposite directions, thereby signaling the higher levels and the Operator-2 network about the fault, so that the alarms are suppressed.

Signaling is achieved through transmission and propagation of AIS protocol data units (PDUs). You must enable AIS explicitly on all the MEPs at the service provider level. A MEP that is configured to issue frames with ETH-AIS information is generally at the server layer and continues to transmit periodic frames with ETH-AIS information until the defect condition is cleared. When a client MEP receives the ETH-AIS frames, it suppresses loss-of-continuity alarms associated with its peer MEPs.



Note that in the absence of AIS, a client MEP resumes generating loss-of-continuity alarms when it detects the loss-of-continuity defect conditions from its server layer.

For point-to-point Ethernet services layer connectivity, a MEP has only one peer MEP. Therefore, there is no ambiguity regarding the peer MEP for which the MEP should suppress alarms when it receives the ETH-AIS information.

For multipoint Ethernet services layer connectivity, a MEP that receives ETH-AIS information cannot determine the exact MEP that encountered the fault and, therefore, cannot isolate the exact peer MEP to suppress the alarms. To avoid this scenario, Y.1731 recommends suppressing alarms for all peer MEPs in the same domain level irrespective of connectivity status in a multipoint Ethernet services layer connectivity setup.

[Table 97 on page 812](#) lists the operational mode commands that you can use in a maintenance domain to check the various parameters pertaining to a MEP.

**Table 97: Operational Mode Commands**

| To Check                                                            | Operational Mode Commands                                                       |
|---------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Whether the AIS configuration is configured correctly on a CFM MEP. | <b>show protocols oam ethernet connectivity-fault-management action-profile</b> |




Table 97: Operational Mode Commands (*continued*)

| To Check                                                       | Operational Mode Commands                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Statistics of AIS frames.                                      | <b>show oam ethernet connectivity-fault-management interfaces detail</b><br><br><b>show oam ethernet connectivity-fault-management mep-statistics</b><br><b>maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i></b><br><b>remote-mep <i>mep-id</i> local-mep <i>mep-id</i></b> |
| Whether any event has occurred that triggered AIS.             | <b>show oam ethernet connectivity-fault-management mep-database</b><br><b>maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i></b><br><b>remote-mep <i>mep-id</i> local-mep <i>mep-id</i></b>                                                                                   |
| Status of CFM sessions for faults that trigger AIS on the MEP. | <b>show oam ethernet connectivity-fault-management interfaces detail</b>                                                                                                                                                                                                                            |

**Terms Defined**

- AIS transmission—A MEP upon detecting a defect condition transmits AIS frames in a direction opposite to its peer MEPs. The periodicity of AIS frames transmission is on the basis of the AIS transmission period. An AIS transmission period of 1 second is recommended. The first AIS frame must always be transmitted immediately following the detection of a defect condition.
- AIS reception—Upon receiving an AIS frame, a MEP examines it to ensure that the frame's maintenance domain level is the same as its own maintenance domain level. The *period* field in the frame indicates the period at which the AIS frames can be expected. When a MEP receives an AIS frame, it detects the defect condition. After detection, when no AIS frames are received within an interval of 3.5 times—the AIS transmission period indicated in the AIS frames received—the MEP clears the AIS defect condition. When the AIS condition is cleared and defects still exist, then the MEPs continue to report alarms.
- AIS PDU format—The fields of the AIS PDU format are:



- MEG Level—Also called the maintenance domain level, it is a 3-bit field that is used to carry the maintenance domain level of the client MEG.
- Version—Value is always 0.

- c. OpCode—Value for this PDU type is AIS (33).
- d. Flags—The first five bits are reserved and are set to 0. The 3-bit information element carried in the three least significant bits are referred to as the period that contains the value of AIS transmission periodicity as illustrated in [Table 98 on page 814](#):

**Table 98: AIS Transmission Periodicity**

| Flags [3:1] | Period Value  | Comments              |
|-------------|---------------|-----------------------|
| 000-011     | Invalid value | Invalid value for AIS |
| 100         | 1s            | 1 frame per second    |
| 101         | Invalid value | Invalid value for AIS |
| 110         | 1 min         | 1 frame per minute    |
| 111         | Invalid value | Invalid value for AIS |

- e. TLV offset—Set to 0.
  - f. End TLV—All-zeroes octet value.
- Server layer and client layer—These layers are part of the ITU-T Recommendation G.805 transport network functional model. This model is based on the concept of layering within a transport network. A transport network is divided into several independent transport layer networks that have a client-server association between adjacent layer networks.
  - Maintenance domain—To enable connectivity fault management (CFM) on an Ethernet interface, maintenance domains, maintenance associations, and maintenance end points (MEPs) are created and configured in a network. You can configure up to eight maintenance domain levels in a network. Each maintenance domain level is a part of the network where the connectivity issues can be monitored and corrected. Provider domain and customer domain are some examples for maintenance domains. Each maintenance domain has a maintenance association. Each maintenance association includes MEPs and maintenance intermediate points (MIPs) in that domain. The MEPs are located at the boundary of the domain and the MIPs are located within the domain. MEPs generate and transmit continuity check messages (CCMs) at configured intervals to the entire maintenance association to check the connectivity in the network.
  - Ethernet services (ETH) layer—A layer in the metro Ethernet network model, where this layer is responsible for the OAM services that are required to support the Ethernet services in the network.

SEE ALSO

*show oam ethernet connectivity-fault-management interfaces*

*show oam ethernet connectivity-fault-management mep-statistics* | 1916

## Ethernet Alarm Indication Signal Overview

ACX Series routers support ITU-T Y.1731 Ethernet Alarm Indication Signal function (ETH-AIS) to provide fault management for service providers. ETH-AIS enables you to suppress alarms when a fault condition is detected. Using ETH-AIS, an administrator can differentiate between faults at customer level or faults at provider level.

The advantage of ETH-AIS is:

- Customers need not raise alarms due to lower level failures.
- Customers can get refund based on service unavailability.

When a fault condition is detected, a maintenance end point (MEP) generates ETH-AIS packets to the configured client levels for a specified duration until the fault condition is cleared. Any MEP configured to generate ETH-AIS packets signals to a level higher than its own. A MEP receiving ETH-AIS recognizes that the fault is at a lower level and then suppresses alarms at current level.

ACX Series routers support ETH-AIS PDU generation for server MEPs based on the following defect conditions:

- Loss of connectivity (physical link loss detection)
- Layer 2 circuit or Layer 2 VPN down

Alarm indication signaling is done through the transmission and propagation of ETH-AIS PDUs. ETH-AIS should be enabled on MEPs. A MEP which is configured to issue packets with ETH-AIS information is generally of server layer and continues to transmit periodic packets with ETH-AIS information until the defect condition is cleared. CFM MEPs, upon receiving ETH-AIS PDUs, suppresses loss of continuity alarms associated with its peer MEPs. A MEP resumes loss of continuity alarm generation upon detecting loss of continuity defect conditions in the absence of an ETH-AIS condition.

For point-to-point Ethernet connectivity, a MEP has only a single peer MEP. Therefore, a MEP suppress alarms on its peer MEP when it receives the ETH-AIS information.

For multi-point Ethernet connectivity, a MEP which receives ETH-AIS information cannot determine the exact MEP encountered a fault condition and therefore it will not be able to isolate the exact peer MEP for alarm suppression. ITU-T Y.1731 recommends suppressing alarms for all peer MEPs irrespective of the connectivity status.

**AIS transmission**—A MEP upon detecting a defect condition transmits ETH-AIS PDUs in a direction opposite to its peer MEPs. The transmission of ETH-AIS PDUs is based on a configured ETH-AIS transmission period. An ETH-AIS transmission period of 1 second is recommended. The first ETH-AIS PDU must be transmitted immediately following the detection of a defect condition.

AIS reception—A MEP upon receiving ETH-AIS PDUs examines it to ensure that its maintenance domain (MD) level corresponds to the same MD level. Upon receiving an ETH-AIS PDU, the MEP detects a defect condition. Following the detection of a defect condition, if there are no ETH-AIS PDUs received within an interval of 3.5 times the ETH-AIS transmission period indicated in the ETH-AIS PDUs received earlier, the MEP clears the defect condition. After the fault condition is cleared, MEPs continue to report alarms.

**NOTE:** ACX Series routers do not support ITU-T Y.1731 ETH-AIS for layer 2 services (bridging).

The following are the limitations for server MEP

- Triggering of ETH-AIS messages over services (Layer 2 circuit and Layer 2 VPN) by the link-loss server MEP is done on a best-effort manner. This is because the transmission of ETH-AIS messages is independent of the service status and there is no guarantee for delivering the ETH-AIS messages before service goes down.
- Pseudowire protection with CFM-MEP session is not monitored by the server-MEP because an entity to monitor pseudowire protection already exists for the service (Layer 2 circuit and Layer 2 VPN).

SEE ALSO

`show oam ethernet connectivity-fault-management mep-statistics`

## Configuring ETH-AIS on a CFM MEP

MX Series routers support ITU-T Y.1731 Ethernet alarm indication signal (ETH-AIS) function to provide fault management for service providers. ETH-AIS enables the service provider to suppress alarms when a fault condition is detected.

The following points are to be noted when ETH-AIS is configured in a maintenance domain:

- Transmitting or receiving of AIS on a MEP does not override the **lowest-priority-defect** statement configured at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]` hierarchy level. Therefore, alarms are reported according to the defect priority configured.
- Alarms are reported even when the higher domain levels exchange CCMs at a faster rate than the lower domain levels.
- Maintenance association intermediate point (MIP) is transparent to ETH-AIS frames—that is, the MIPs do not perform any action in response to ETH-AIS frames.
- When the service provider requests the MEP to generate an AIS for a lower level or for the same level, the request is rejected.

- AIS generation is stopped when the MEP clears the remote MEP within the maintenance association.
- When the **auto-discovery** statement is enabled for a MEP, the remote MEP information is cleared after the configured hold interval expires.

The following tasks explain how to enable ETH-AIS in a maintenance domain, configure an action to be taken when a defect is detected, and to attach the action profile to a CFM MEP:

1. [Configuring an Action Profile | 817](#)
2. [Configuring an Action to Be Taken When an AIS Alarm Is Detected | 818](#)
3. [Attaching the Action Profile to a CFM MEP | 819](#)

### **Configuring an Action Profile**

To configure an action profile for ETH-AIS:

1. Go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management
```

2. Configure an action profile to use when one or more remote MEPs are down.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# edit action-profile action-profile-name
```

3. Configure an event that needs to be monitored.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name]
user@host# edit event
```

4. Configure the defect condition that generates an alarm indication signal.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event]
user@host# edit ais-trigger-condition
```

5. Configure the **adjacency-loss** statement to inform the operator when the physical connectivity is lost between the peer MEPs.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event
ais-trigger-condition]
```

```
user@host# set adjacency-loss
```

6. Configure the **all-defects** statement to inform the operator that all possible defects must be considered to raise the alarm indication signal.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event
  ais-trigger-condition]
user@host# set all-defects
```

7. Configure the **cross-connect-ccm** statement to inform the operator when cross-connect continuity check messages (CCMs) are received by the MEP and to raise an alarm indication signal in response.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event
  ais-trigger-condition]
user@host# set cross-connect-ccm
```

8. Configure the **erroneous-ccm** statement to inform the operator when CCMs with unexpected MEP ID or maintenance domain level are received by the MEP and an AIS alarm is raised in response.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event
  ais-trigger-condition]
user@host# set erroneous-ccm
```

9. Configure the **receive-ais** statement to inform the operator that an AIS message has been received from the peer MEP in its own maintenance level.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event
  ais-trigger-condition]
user@host# set receive-ais
```

### ***Configuring an Action to Be Taken When an AIS Alarm Is Detected***

Configure an action to be taken when an AIS alarm is detected.

1. Go to the `[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action]` hierarchy level.

```
[edit]
```

```
user@host# edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name
action
```

2. Configure the **log-and-generate-ais** statement to log the event that generated the AIS message.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action]
user@host# edit log-and-generate-ais
```

3. Configure the interval between AIS messages that are to be received by the MEP as 1 minute or 1 second.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action
log-and-generate-ais]
user@host# set interval (1m | 1s)
```

4. Configure the server maintenance domain level range of the MEP from 1 through 7.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action
log-and-generate-ais]
user@host# set level level
```

5. Configure the 802.1p priority of the AIS packet from 1 through 7.

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action
log-and-generate-ais]
user@host# set priority level
```

### **Attaching the Action Profile to a CFM MEP**

After configuring an event and an action to be monitored in an action profile, you must attach the action profile to a CFM MEP.

1. Go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management
```

2. Configure the maintenance domain with a name.

```
[edit protocols oam ethernet connectivity-fault-management]
user@host# edit maintenance-domain md-name
```

3. Configure the maintenance domain with a client maintenance entity group (MEG) level or maintenance association level—the level which the client layer maintenance association intermediate point (MIPs) and the MEPs exist—from 0 through 7.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name]
user@host# edit level level
```

**NOTE:** You cannot configure a maintenance domain level that is lower than or equal to the maintenance association level that it is associated with.

4. Configure the maintenance association.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name]
user@host# edit maintenance-association ma-name
```

5. Configure the continuity check that is performed on all the MEPs in a domain level by sending CCMs with an interval between two CCMs—100 milliseconds, 10 milliseconds, 1 second, 10 seconds, 1 minute, or 10 minutes—and the number of CCMs that are to be lost before marking a MEP as down.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name]
user@host# set continuity-check interval (100ms | 10m | 10ms | 1m | 1s)
user@host# set continuity-check loss-threshold value
```

6. Configure the MEP with an identifier from 1 through 8192.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name]
user@host# set mep mep-id
```

7. Attach the configured action profile to the MEP.



```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id]
user@host# set action-profile action-profile-name
```

8. Configure the interface of the MEP over which the CCMs are transmitted.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id]
user@host# set interface interface-name
```

9. Configure the direction for the CCMs to travel to the next MEP as up or down.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id]
user@host# set direction (down | up)
```

10. Configure the 802.1p priority for the CCMs and link-trace packet from 0 through 7.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id]
user@host# set priority priority-value
```

#### SEE ALSO

[ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)  
[Ethernet Alarm Indication Signal Overview | 815](#)

### Configuring Alarm Indication Signal on ACX Series Routers

ACX Series routers support ITU-T Y.1731 Ethernet Alarm Indication Signal function (ETH-AIS) to provide fault management for service providers. ETH-AIS enables you to suppress alarms when a fault condition is detected.

To support ETH-AIS transmission, the following configuration information is required by a CFM MEP:

- Client Maintenance Entity Group level—Maintenance Entity Group (MEG) level at which the immediate client layer Maintenance Domain Intermediate Points (MIPs) and Maintenance Association End Points (MEPs) exist.

- ETH-AIS transmission period—Determines the ETH-AIS PDU transmission interval.
- Priority—Determines the priority of packets with ETH-AIS information. This is optional.

To configure ETH-AIS in CFM MEP, you need to:

- Configure an action profile with ETH-AIS action
- Attach the action profile to the CFM MEP

To configure an action profile with ETH-AIS action, include the following statements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```
[edit protocols oam ethernet connectivity-fault-management]
action-profile action-profile-name {
  event {
    adjacency-loss;
    all-defects;
    cross-connect-ccm;
    errored-ccm;
    receive-ais;
  }
  action {
    log-and-generate-ais {
      level [1-7];
      interval 1s | 1m ;
      priority [0-7];
    }
    log-ais;
  }
}
```

To attach an action profile to a CFM MEP, include the following statements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```
maintenance-domain maintenance-domain-name {
  level level-number;
  maintenance-association maintenance-domain-name {
    continuity-check {
      interval 1s;
      loss-threshold 3;
    }
    mep mep-id {
      interface interface-name;
      direction up | down;
      priority priority-value;
    }
  }
}
```

```

        action-profile action-profile-name;
    }
}

```

**NOTE:** You cannot configure a maintenance domain level that is lower than or equal to the level that it is associated with.

To support ETH-AIS transmission, the following configuration information required by a server MEP:

- Server MEP definition—Defines the association of server MEP identifier to the server layer.
  - For Layer 2 circuit and Layer 2 VPN, the logical interface connected to a customer network (UNI) would be the identifier for the server layer that needs to be monitored by the server MEP.
  - For physical link loss detection, the physical interface under Ethernet protocol would be the identifier for the server layer that needs to be monitored by the server MEP.
- Association of server MEP defect—Defines the association of server MEP defects to ETH-AIS action.
- Association action profile and server MEP—Defines the binding of server MEP and action profile.

To configure ETH-AIS in server MEP, you need to:

- Create an action profile with ETH-AIS action for server MEP defects.
- Attach the action profile to a server MEP

To create an action profile, include the following statements at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```

[edit protocols oam ethernet connectivity-fault-management]
action-profile action-profile-name {
    event {
        server-mep-defects {
            link-loss-defect;
            l2circuit-defect;
            l2vpn-defect;
        }
    }
    action {
        log-and-generate-ais {
            level 1...n;
            interval 1 second | 1 minute;
        }
    }
}

```

```

        priority dot1p [range 0-7];
    }
}

```

To attach an action profile to a server MEP, include the following statement at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level:

```

[edit protocols oam ethernet connectivity-fault-management]
server-mep mep-identifier {
    protocol l2circuit | l2vpn | ethernet {
        interface interface-name;
    }
    action-profile action-profile-name;
}

```

#### SEE ALSO

[ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)

[Ethernet Alarm Indication Signal Overview | 815](#)

#### RELATED DOCUMENTATION

[ITU-T Y.1731 Ethernet Service OAM Overview | 689](#)

[Configuring Ethernet Frame Delay Measurement Sessions | 709](#)

[Configuring Ethernet Frame Loss Measurement | 747](#)

[Configuring Connectivity Fault Management \(CFM\) | 553](#)

## Inline Transmission Mode

#### IN THIS SECTION

- [Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling | 825](#)
- [Enabling Inline Transmission of Link Fault Management Keepalives for Maximum Scaling | 826](#)
- [Enabling Inline Mode Of Performance Monitoring To Achieve Maximum Scaling | 830](#)

- Supported Inline CCM and Inline PM Scaling Values | 833
- Enabling the Hardware-Assisted Timestamping Option | 835

Use this topic to understand what inline transmission is and how to enable it for maximum scaling for CFM, LFM, and performance monitoring functions.

## Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling

Scaling is the ability of a system to handle increasing amounts of work and to continue to function well. Scaling can refer to increasing capacity and the ability to handle increasing workload, number of subscribers or sessions, hardware components, and so on. Continuity check protocol is used for fault detection within a maintenance association. The maintenance association end points (MEPs) send continuity check messages (CCMs) periodically. The time between the transmissions of CCMs is known as the interval. The receiving MEP maintains a database of all MEPs in the maintenance association.

By default, CCMs are transmitted by the CPU of a line card, such as a Modular Port Concentrator (MPC). If the duration between transmissions of CCMs is low or if the CCMs for a specific line card scale, then we recommend that you delegate transmission of CCMs to the forwarding ASIC (that is, to the hardware) by enabling inline transmission of CCMs. Inline transmission of CCMs is also known as inline keepalives or Inline-KA. Inline transmission enables the system to handle more connectivity fault management (CFM) sessions per line card. By enabling inline transmission of CCMs, you can achieve maximum scaling of CCMs.

To enable inline transmission of CCMs, perform the following steps:

1. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level.

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management performance-monitoring
```

2. Delegate transmission of CCMs to hardware by enabling hardware-assisted keepalives.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set hardware-assisted-keepalives enable
```

**NOTE:** Inline transmission of CCMs is not enabled when there is a CFM session already established. To enable inline transmission, you must first deactivate the CFM session using the **deactivate** command and then reactivate the CFM session using the **activate** command.

To disable inline transmission, use the **hardware-assisted-keepalives disable** statement. After disabling inline transmission, you must reboot the router for the changes to take effect.

#### SEE ALSO

[Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades | 599](#)

### Enabling Inline Transmission of Link Fault Management Keepalives for Maximum Scaling

Scaling is the ability of a system to handle increasing amounts of work and to continue to function well. Scaling can refer to increasing capacity and the ability to handle increasing workload, number of subscribers or sessions, hardware components, and so on.

By default, LFM keepalive packets are transmitted by the periodic packet management **ppm** process on the line-card. You can delegate transmission of LFM keepalive packets to the forwarding ASIC (that is, to the hardware) by enabling inline transmission. Inline transmission of LFM keepalives is also known as inline keepalives or Inline-KA. By enabling inline transmission of LFM keepalive packets, you can achieve maximum scaling of keepalive packets, reduction of the load on the **ppm** process, and support LFM in-service software upgrade (ISSU) for non-juniper peers (for a keepalive interval of 1 second).

**NOTE:** Do not enable or disable inline transmission of LFM when an LFM session is already established. To enable or disable inline transmission, you must first deactivate the existing established LFM session using the **deactivate** command, and then reactivate the LFM session using the **activate** command after enabling or disabling inline LFM.

Before you enable inline transmission of LFM keepalive packets, complete the following tasks:

- Verify if any LFM session is online and active. To verify if any existing or established LFM session is online and active, issue the following command:

```
user@host> show oam ethernet link-fault-management detail
```

```

Oct 18 02:04:17
Interface: ge-0/0/0
  Status: Running, Discovery state: Active Send Local
  Transmit interval: 1000ms, PDU threshold: 3 frames, Hold time: 0ms
  Peer address: 00:00:00:00:00:00
  Flags:0x8
  OAM receive statistics:
    Information: 0, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 0, Organization specific: 0
  OAM flags receive statistics:
    Critical event: 0, Dying gasp: 0, Link fault: 0
  OAM transmit statistics:
    Information: 28, Event: 0, Variable request: 0, Variable response: 0 =
after waiting for a while count increased by 15
    Loopback control: 0, Organization specific: 0
  OAM received symbol error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame period error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame seconds error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM transmitted symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM transmitted frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  Loopback tracking: Disabled, Loop status: Unknown
  Detect LOC: Disabled, LOC status: Unknown

```

The OAM transmit statistics reflect that the **ppm** process is handling the transmission of LFM keepalive packets.

- Deactivate the LFM session so that you can enable inline LFM mode. To deactivate the LFM session, issue the following command:

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

- Commit the configuration. To commit the configuration, issue the following command:

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

To enable inline transmission of LFM keepalive packets, perform the following steps:

1. In configuration mode, go to the **[edit protocols oam ethernet link-fault-management]** hierarchy level.

```
[edit]
user@host# edit protocols oam ethernet link-fault-management
```

2. Delegate transmission of LFM keepalive packets to hardware by enabling hardware-assisted keepalives.

```
[edit protocols oam ethernet link-fault-management]
user@host# set hardware-assisted-keepalives
```

3. Commit the configuration.

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

4. Re-activate the LFM session as follows:

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

5. Commit the configuration.

```
[edit]
```



```
user@host # commit
```

6. Verify that the transmission of LFM keepalive packets is delegated from the **ppm** process to the hardware. To verify that you have enabled inline transmission, issue the following command:

```
user@host> show oam ethernet link-fault-management detail
```

```
Oct 18 02:05:05
Interface: ge-0/0/0
  Status: Running, Discovery state: Active Send Local
  Transmit interval: 1000ms, PDU threshold: 3 frames, Hold time: 0ms
  Peer address: 00:00:00:00:00:00
  Flags:0x8
  OAM receive statistics:
    Information: 0, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 0, Organization specific: 0
  OAM flags receive statistics:
    Critical event: 0, Dying gasp: 0, Link fault: 0
  OAM transmit statistics:
    Information: 1, Event: 0, Variable request: 0, Variable response: 0 =
even after 10 seconds count is still 1
    Loopback control: 0, Organization specific: 0
  OAM received symbol error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame period error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame seconds error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM transmitted symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM transmitted frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current frame error event information:
```

```

Events: 0, Window: 0, Threshold: 1
Errors in period: 0, Total errors: 0
Loopback tracking: Disabled, Loop status: Unknown
Detect LOC: Disabled, LOC status: Unknown

```

The OAM transmit statistics are not updated. When you enable inline transmission of LFM keepalive packets, the OAM transmit statistics are not updated.

To disable inline LFM, verify if any existing established LFM session is online and active. Deactivate the LFM session and commit. Disable inline LFM by deleting the **hardware-assisted-keepalives** statement and commit. Then, reactivate LFM session and commit the configuration.

## SEE ALSO

*hardware-assisted-keepalives (lfm)*

## Enabling Inline Mode Of Performance Monitoring To Achieve Maximum Scaling

Performance monitoring is useful for studying the traffic pattern in a network over a period of time. It helps to identify network problems before you are impacted by network defects.

By default, performance monitoring packets are handled by the CPU of a line-card, such as Modular Port Concentrator (MPC). Enabling inline mode of performance monitoring delegates the processing of the protocol data units (PDUs) to the forwarding ASIC (that is, to the hardware). By enabling inline mode of performance monitoring, the load on the CPU of the line-card is reduced and you can configure an increased number of performance monitoring sessions and achieve maximum scaling for service OAM performance monitoring sessions. On MX Series routers, you can configure inline mode of performance monitoring only if the network services mode on the router is configured to **enhanced-ip** and enhanced connectivity fault management (**enhanced-cfm-mode**) is configured.

By enabling inline mode of performance monitoring, you can achieve maximum scaling for performance monitoring sessions. To achieve maximum scaling for performance monitoring sessions, you must enable scaling of continuity check messages (CCMs) sessions. To enable scaling of CCM sessions, enable inline transmission of continuity check messages. For more information on inline transmission of continuity check messages, see [“Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling” on page 825](#). To view the supported scaling values for CCM and PM, see [“Supported Inline CCM and Inline PM Scaling Values” on page 833](#).

Inline mode of performance monitoring is supported only for proactive mode of frame delay measurement (Two-way Delay Measurements) and synthetic loss measurements (SLM) sessions. Performance monitoring functions configured using the iterator profile (CFM) are referred to as proactive performance monitoring.

Inline mode of performance monitoring for frame loss measurement using service frames (LM) is not supported.

**NOTE:** MPC3E (MX-MPC3E-3D) and MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE) do not support inline mode of performance monitoring. User-defined Data TLV is not supported if you have configured inline mode of performance monitoring. Also, only 12 history records per PM sessions are supported.

We recommend that you enable inline mode of performance monitoring before you configure the performance monitoring sessions as the change may interfere with the existing performance monitoring sessions.

To enable inline mode of performance monitoring, perform the following steps:

1. In configuration mode, go to the **[edit chassis]** hierarchy level and configure the network services mode of the router. The network service mode of the router must be configured as **enhanced ip** to enable enhanced connectivity fault management (CFM) mode.

**NOTE:** If the network services mode is not **enhanced-ip**, and you have enabled enhanced CFM, the following warning message is displayed:

**[edit protocols oam ethernet]**

**'connectivity-fault-management'**

**enhanced ip is not effective please configure enhanced ip and give router reboot**

**[edit chassis]**

**user@host# set network-services enhanced-ip**

2. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level and enable enhanced connectivity fault management mode by using the **enhanced-cfm-mode** option.

**[edit]**

**user@host# set protocols oam ethernet connectivity-fault-management **enhanced-cfm-mode****

3. In configuration mode, go to the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level. Configure the enhanced iterator profile by using the **enhanced-sla-iterator** option and specify the measurement interval by using the **measurement-interval** option.

```
[edit]
user@host# edit protocols oam ethernet connectivity-fault-management performance-monitoring
enhanced-sla-iterator measurement-interval value
```

4. Enable inline performance monitoring.

**NOTE:** You can enable inline mode of performance monitoring for both the originator and the responder of the service OAM performance monitoring sessions by using the **hardware-assisted-pm** command.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set hardware-assisted-pm
```

5. (Optional) Enable inline transmission of CCMs to enable better scaling if inline transmission of CCMs is not automatically enabled.

**NOTE:** You can achieve better scaling if both inline performance monitoring and inline transmission of CCMs is enabled.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# set hardware-assisted-keepalives enable
```

6. Commit the configuration.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@host# commit
```

## SEE ALSO

[Enabling Enhanced Connectivity Fault Management Mode | 596](#)

[Network Services Mode Overview](#)

[hardware-assisted-pm](#)

## Supported Inline CCM and Inline PM Scaling Values

This topic lists the scaling values for inline mode of performance monitoring and inline transmission of continuity check messages. The scaling values are based on the different cycle-time interval values. Each table lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure inline CCM, enhanced CFM, and enhanced PM by using the **hardware-assisted-keepalives**, **enhanced-cfm-mode**, and **hardware-assisted-pm** options.

**NOTE:** The scaling values do not consider the load from other protocols in the system and so the actual realized scaling values for line card and chassis vary depending on other protocol configurations and scaling in the system. We recommend that you configure DDoS for CFM. Limit the number of CFM packets, that are sent to the CPU of the line card, to 3000. Limiting the number of packets safeguards the CPU from scaled CFM configurations of various CFM protocol events.

[Table 99 on page 833](#) lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure both the CCM interval and the PM interval as 1 second.

**Table 99: Scaling Values for CFM and PM (CCM Interval: 1 sec and PM Interval: 1 sec )**

| CFM Line Card Scale | PM Line Card Scale | CFM Chassis Scale | PM Chassis Scale |
|---------------------|--------------------|-------------------|------------------|
| 4000                | 4500               | 16000             | 16000            |
| 6000                | 3750               | 16000             | 16000            |
| 7000                | 3375               | 16000             | 16000            |
| 8000                | 3000               | 16000             | 16000            |

[Table 100 on page 833](#) lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure the CCM interval as 1 second and the PM interval as 100 milliseconds.

**Table 100: Scaling Values for CFM and PM (CCM Interval: 1 sec and PM interval: 100 ms )**

| CFM Line Card Scale | PM Line Card Scale | CFM Chassis Scale | PM Chassis Scale |
|---------------------|--------------------|-------------------|------------------|
| 4000                | 450                | 12000             | 4000             |
| 6000                | 375                | 12000             | 4000             |

**Table 100: Scaling Values for CFM and PM (CCM Interval: 1 sec and PM interval: 100 ms ) (continued)**

| CFM Line Card Scale | PM Line Card Scale | CFM Chassis Scale | PM Chassis Scale |
|---------------------|--------------------|-------------------|------------------|
| 7000                | 337                | 12000             | 4000             |
| 8000                | 300                | 12000             | 4000             |

[Table 101 on page 834](#) lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure the CCM interval as 100 milliseconds and the PM interval as 1 second.

**Table 101: Scaling Values for CFM and PM (CCM Interval: 100 ms and PM interval: 1 sec )**

| CFM Line Card Scale | PM Line Card Scale | CFM Chassis Scale | PM Chassis Scale |
|---------------------|--------------------|-------------------|------------------|
| 4000                | 3000               | 8000              | 6000             |
| 3000                | 3750               | 8000              | 6000             |
| 2000                | 4500               | 8000              | 6000             |
| 1000                | 4500               | 8000              | 6000             |

[Table 102 on page 834](#) lists the maximum number of connectivity fault management (CFM) sessions and performance monitoring (PM) sessions per line card and per chassis when you configure both the CCM interval and the PM interval as 100 milliseconds.

**Table 102: Scaling Values for CFM and PM (CCM Interval: 100 ms and PM interval: 100 ms )**

| CFM Line Card Scale | PM Line Card Scale | CFM Chassis Scale | PM Chassis Scale |
|---------------------|--------------------|-------------------|------------------|
| 4000                | 300                | 8000              | 3000             |
| 3000                | 375                | 8000              | 3000             |
| 2000                | 450                | 8000              | 3000             |
| 1000                | 450                | 8000              | 3000             |

SEE ALSO

[enhanced-cfm-mode](#) | [1023](#)

[hardware-assisted-pm](#)

## Enabling the Hardware-Assisted Timestamping Option

By default, Ethernet frame delay measurement uses software for timestamping transmitted and received ETH-DM frames. For Ethernet interfaces, you can optionally use hardware timing to assist in the timestamping of received ETH-DM frames to increase the accuracy of delay measurements.

Enabling hardware-assisted timestamping of received frames can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction.

To enable Ethernet frame delay measurement hardware assistance on the reception path, include the **hardware-assisted-timestamping** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      performance-monitoring {
        hardware-assisted-timestamping;
      }
    }
  }
}
```

### SEE ALSO

[hardware-assisted-timestamping](#) | 1028

### RELATED DOCUMENTATION

[Introduction to OAM Connectivity Fault Management \(CFM\)](#) | 546

[Configuring Continuity Check Messages](#) | 648

[Introduction to OAM Link Fault Management \(LFM\)](#) | 655

# 4

PART

## Troubleshooting Information

---

[Monitoring and Troubleshooting Ethernet Interfaces](#) | **839**

---





# Monitoring and Troubleshooting Ethernet Interfaces

## IN THIS CHAPTER

- [Passive Monitoring on Ethernet Interfaces Overview | 839](#)
- [Enabling Passive Monitoring on Ethernet Interfaces | 841](#)
- [Link Degradate Monitoring Overview | 843](#)
- [Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 846](#)
- [Performing Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces | 859](#)
- [Performing Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces | 878](#)
- [Configuring Interface Diagnostics Tools to Test the Physical Layer Connections | 895](#)
- [Locating the Fast Ethernet and Gigabit Ethernet LINK Alarm and Counters | 902](#)
- [Troubleshooting: 10-Gigabit Ethernet Port Stuck in Down State | 908](#)
- [Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test | 910](#)

## 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 839](#). 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 | 839](#)

*Configuring Multiservice Physical Interface Properties*

*Junos OS Services Interfaces Library for Routing Devices*

*Ethernet Interfaces User Guide for Routing Devices*

## Link Degrad Monitoring Overview

Link degrade monitoring enables you to monitor the quality of physical links on Ethernet interfaces (10-Gigabit, 40-Gigabit, and 100-Gigabit) 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. When configured 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. 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.

**NOTE:** Layer 2 and Layer 3 protocols already support physical link monitoring. So do Ethernet links through the Link Fault Signaling (LFS) protocol. However, these existing mechanisms cannot detect BER ranges that are very low (for example,  $10^{-13}$  through  $10^{-5}$ ).

**Supported Platforms**

[Table 103 on page 845](#) lists the platform series and line cards that support link degrade monitoring.

Table 103: Line Cards that Support Link Degrade 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> <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 degrade monitoring is not supported on MACsec-enabled MPC7E-10G and MIC-MACSEC-MRATE.</p> | <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>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> |

## RELATED DOCUMENTATION

[Physical Interface Damping Overview](#)
[Fast Reroute Overview](#)
[link-degrade-monitor | 1227](#)
[thresholds | 1384](#)



recovery | 1335

request interface link-degrade-recover | 1476

## Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

### IN THIS SECTION

- Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces | 846
- Monitor Fast Ethernet and Gigabit Ethernet Interfaces | 847
- Fiber-Optic Ethernet Interface Specifications | 857

### Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

#### Purpose

To monitor Fast Ethernet and Gigabit Ethernet interfaces and begin the process of isolating interface problems when they occur.

#### Action

Table 104 on page 846 provides links and commands for monitoring Fast Ethernet and Gigabit Ethernet interfaces.

Table 104: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces

| Tasks                                                                                                                          | Command or Action                                                          |
|--------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| <b>"Monitor Fast Ethernet and Gigabit Ethernet Interfaces" on page 847</b>                                                     |                                                                            |
| 1. <a href="#">Display the Status of Fast Ethernet Interfaces on page 847</a>                                                  | <code>show interfaces terse (fe*   ge*)</code>                             |
| 2. <a href="#">Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface on page 850</a>                    | <code>show interfaces (fe-fpc/pic/port   ge-fpc/pic/port)</code>           |
| 3. <a href="#">Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface on page 852</a> | <code>show interfaces (fe-fpc/pic/port   ge-fpc/pic/port) extensive</code> |
| 4. <a href="#">Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface</a>                                        | <code>monitor interface (fe-fpc/pic/port   ge-fpc/pic/port)</code>         |

Table 104: Checklist for Monitoring Fast Ethernet and Gigabit Ethernet Interfaces *(continued)*

| Tasks                                                                        | Command or Action |
|------------------------------------------------------------------------------|-------------------|
| 5. <a href="#">Fiber-Optic Ethernet Interface Specifications on page 857</a> |                   |

**Meaning**

You can use the above described commands to monitor and to display the configurations for Fast Ethernet and Gigabit Ethernet interfaces.

SEE ALSO

|                                                                         |
|-------------------------------------------------------------------------|
| <a href="#">Display the Status of Gigabit Ethernet Interfaces   849</a> |
| <a href="#">Display the Status of Fast Ethernet Interfaces   847</a>    |

**Monitor Fast Ethernet and Gigabit Ethernet Interfaces**

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:

- 1. [Display the Status of Fast Ethernet Interfaces | 847](#)
- 2. [Display the Status of Gigabit Ethernet Interfaces | 849](#)
- 3. [Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface | 850](#)
- 4. [Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface | 852](#)
- 5. [Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 855](#)

***Display the Status of Fast Ethernet Interfaces***

**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

```
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 105 on page 848](#).

**Table 105: Status of Fast Ethernet Interfaces**

| Physical Interface                                     | Logical Interface                                        | Status Description                                                                                                                                                                                                  |
|--------------------------------------------------------|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>fe-2/1/0</b><br><b>Admin Up</b><br><b>Link Up</b>   | <b>fe-2/1/0.0</b><br><b>Admin Up</b><br><b>Link Up</b>   | This interface has both the physical and logical links up and running.                                                                                                                                              |
| <b>fe-3/0/2</b><br><b>Admin Up</b><br><b>Link Down</b> | <b>fe-3/0/2.0</b><br><b>Admin Up</b><br><b>Link Down</b> | 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.                                                                       |
| <b>fe-4/1/0</b><br><b>Admin Down</b><br><b>Link Up</b> | <b>fe-4/1/0.0</b><br><b>Admin Up</b><br><b>Link Down</b> | 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. |

Table 105: Status of Fast Ethernet Interfaces *(continued)*

| Physical Interface | Logical Interface | Status Description                                           |
|--------------------|-------------------|--------------------------------------------------------------|
| fe-4/1/2           | fe-4/1/2.0        | This interface has both the physical and logical links down. |
| Admin Up           | Admin Up          |                                                              |
| Link Down          | Link Down         |                                                              |

SEE ALSO

| [Display the Status of Gigabit Ethernet Interfaces](#) | 849

*Display the Status of Gigabit Ethernet Interfaces*

**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 106 on page 850](#) for a description of what the output means.

Table 106: 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

[Display the Status of Fast Ethernet Interfaces | 847](#)

### *Display the Status of a Specific Fast Ethernet or Gigabit Ethernet Interface*

#### Purpose

To display the status of a specific Fast Ethernet or Gigabit Ethernet interface when you need to investigate its status further, use the following Junos OS CLI operational mode command:

#### Action

```
user@host> show interfaces (fe-fpc/pic/port | ge-fpc/pic/port)
```

#### Sample Output 1

The following sample output is for a Fast Ethernet interface with the physical link up:

```
user@host> show interfaces fe-2/1/0
Physical interface: fe-2/1/0, Enabled, Physical link is Up
  Interface index: 31, SNMP ifIndex: 35
  Description: customer connection
```

```

Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
Speed: 100mbps, Loopback: Disabled, Flow control: Enabled
Device flags      : Present Running
Interface flags: SNMP-Traps
Link flags        : None
Current address: 00:90:69:86:71:1b, Hardware address: 00:90:69:86:71:1b
Input rate       : 25768 bps (11 pps), Output rate: 1576 bps (3 pps)
Active alarms    : None
Active defects   : None
Logical interface fe-2/1/0.0 (Index 2) (SNMP ifIndex 43)
  Flags: SNMP-Traps, Encapsulation: ENET2
  Protocol inet, MTU: 1500, Flags: Is-Primary
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.116.151.218/29, Local: 10.119.115.217
    Broadcast: 10.116.151.225

```

## Sample Output 2

The following output is for a Gigabit Ethernet interface with the physical link up:

```

user@host> show interfaces ge-3/1/0
Physical interface: ge-3/1/0, Enabled, Physical link is Up
  Interface index: 41, SNMP ifIndex: 55
  Description: customer connection
  Link-level type: Ethernet, MTU: 1514, Source filtering: Disabled
  Speed: 1000mbps, Loopback: Disabled, Flow control: Enabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps
  Link flags        : None
  Current address: 00:90:69:85:71:99, Hardware address: 00:90:69:85:71:99
  Input rate       : 7412216 bps (1614 pps), Output rate: 2431184 bps (1776 pps)
  Active alarms    : None
  Active defects   : None
  Logical interface ge-3/1/0.0 (Index 11) (SNMP ifIndex 57)
    Flags: SNMP-Traps, Encapsulation: ENET2
    Protocol inet, MTU: 1500
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.117.65.192/30, Local: 10.115.65.193
      Broadcast: 10.115.65.195

```

## Meaning

The first line of sample output 1 and 2 shows that the physical link is up. This means that the physical link is healthy and can pass packets. Further down the sample output, look for active alarms and defects. If you see active alarms or defects, to further diagnose the problem, see Step 3, [“Display Extensive Status Information for a Specific Fast Ethernet or Gigabit Ethernet Interface” on page 852](#), 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***

#### **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:
    Total octets          439703575      Transmit 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 107 on page 854](#) for a description of errors to look for.

**Table 107: 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” on page 902](#) for an explanation of Fast Ethernet and Gigabit Ethernet alarms.

[Table 108 on page 854](#) lists and describes some MAC statistics errors to look for.

**Table 108: 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.                                                                                                                                                                                                                                                                                 |
| <b>MAC pause frames</b>   | The number of MAC control frames with <b>pause</b> operational code.                                                                                                                                                                                                                                              |

Table 108: MAC Statistics Errors (*continued*)

| Error                  | Meaning                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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 an alignment error.</p> <p>Note that it is entirely normal for fragment frames to increment because both runs (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 109 on page 855](#) explains the autonegotiation information of the **show interface *interface-name* extensive** command output.

Table 109: 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.                                                                     |
| <b>Negotiation status: Complete<br/>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:<br/>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

#### 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” on page 857](#) for the fiber-optic Ethernet interface specifications.
3. Measure the transmit light level on the Tx port to verify that it is within specification. See [“Fiber-Optic Ethernet Interface Specifications” on page 857](#) for the optical specifications.

## Fiber-Optic Ethernet Interface Specifications

[Table 110 on page 857](#) shows the specifications for fiber-optic interfaces for Juniper Networks routers.

**Table 110: 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 110: 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\*](#)


---

[\*Understanding 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 | 859](#)
- [Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface | 860](#)
- [Create a Loopback | 861](#)
- [Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 864](#)
- [Configure a Static Address Resolution Protocol Table Entry | 868](#)
- [Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 873](#)
- [Ping the Fast Ethernet or Gigabit Ethernet Interface | 874](#)
- [Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 875](#)
- [Diagnose a Suspected Circuit Problem | 877](#)

## Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

Purpose

To use loopback testing to isolate Fast Ethernet and Gigabit Ethernet interface problems.

Action

[Table 111 on page 859](#) provides links and commands for using loopback testing for Fast Ethernet and Gigabit Ethernet interfaces.

Table 111: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces

| Tasks                                                                                                                  | Command or Action                                                                     |
|------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <a href="#">“Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface” on page 860</a> |                                                                                       |
| 1. <a href="#">Create a Loopback on page 861</a>                                                                       |                                                                                       |
| a. <a href="#">Create a Physical Loopback for a Fiber-Optic Interface on page 861</a>                                  | Connect the transmit port to the receive port.                                        |
| b. <a href="#">Create a Loopback Plug for an RJ-45 Ethernet Interface on page 862</a>                                  | Cross pin 1 (TX+) and pin 3 (RX+) together, and pin 2 (TX-) and pin 6 (RX-) together. |

**Table 111: Checklist for Using Loopback Testing for Fast Ethernet and Gigabit Ethernet Interfaces** (*continued*)

| Tasks                                                                                                 | Command or Action                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| c. <a href="#">Configure a Local Loopback on page 863</a>                                             | <code>[edit interfaces <i>interface-name</i><br/>(fastether-options   gigether-options)]<br/>set loopback<br/>show<br/>commit</code>                                                                                                                             |
| 2. <a href="#">Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 864</a>      | <code>show interfaces (fe-fpc/pic/port   ge-fpc/pic/port)</code>                                                                                                                                                                                                 |
| 3. <a href="#">Configure a Static Address Resolution Protocol Table Entry on page 868</a>             | <code>show interfaces ge-fpc/pic/port<br/>[edit interfaces <i>interface-name</i> unit<br/><i>logical-unit-number</i> family inet address <i>address</i>]<br/>set arp <i>ip-address</i> mac <i>mac-address</i> show<br/>commit<br/>run show arp no-resolve</code> |
| 4. <a href="#">Clear Fast Ethernet or Gigabit Ethernet Interface Statistics on page 873</a>           | <code>clear interfaces statistics fe-fpc/pic/port  <br/>ge-fpc/pic/port</code>                                                                                                                                                                                   |
| 5. <a href="#">Ping the Fast Ethernet or Gigabit Ethernet Interface on page 874</a>                   | <code>ping <i>remote-IP-address</i> bypass-routing interface<br/>(fe-fpc/pic/port   ge-fpc/pic/port count 100<br/>rapid</code>                                                                                                                                   |
| 6. <a href="#">Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics on page 875</a> | <code>show interfaces (fe-fpc/pic/port   ge-fpc/pic/port<br/>) extensive</code>                                                                                                                                                                                  |
| <a href="#">“Diagnose a Suspected Circuit Problem” on page 877</a>                                    | Perform Steps 2 through 8 from <a href="#">“Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface” on page 860</a> .                                                                                                          |

## Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface

### Problem

**Description:** When you suspect a hardware problem, take the following steps to help verify if there is a problem.

### Solution

To diagnose a suspected hardware problem with the Ethernet interface, follow these steps:

- [Create a Loopback on page 861](#)
- [Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up on page 864](#)
- [Configure a Static Address Resolution Protocol Table Entry on page 868](#)
- [Clear Fast Ethernet or Gigabit Ethernet Interface Statistics on page 873](#)
- [Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics on page 875](#)

## Create a Loopback

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).

1. [Create a Physical Loopback for a Fiber-Optic Interface | 861](#)
2. [Create a Loopback Plug for an RJ-45 Ethernet Interface | 862](#)
3. [Configure a Local Loopback | 863](#)

### *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 Loopback Plug for an RJ-45 Ethernet Interface | 862](#)

[Configure a Local Loopback | 863](#)

---



---

[Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 864](#)


---

[Configure a Static Address Resolution Protocol Table Entry | 868](#)


---

### 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 27 on page 862](#) illustrates how to create a loopback plug for an RJ-45 Ethernet interface.

**Figure 27: 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

---

[Configure a Local Loopback | 863](#)


---

[Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 864](#)


---

[Configure a Static Address Resolution Protocol Table Entry | 868](#)


---

## 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 | gigether-options)]** hierarchy level, places the interface into local loopback mode.

**NOTE:** Remember to delete the loopback statement after completing the test.

## SEE ALSO

[Create a Loopback Plug for an RJ-45 Ethernet Interface | 862](#)  
[Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 864](#)  
[Configure a Static Address Resolution Protocol Table Entry | 868](#)

## SEE ALSO

[Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up | 864](#)  
[Configure a Static Address Resolution Protocol Table Entry | 868](#)  
[Clear Fast Ethernet or Gigabit Ethernet Interface Statistics | 873](#)  
[Ping the Fast Ethernet or Gigabit Ethernet Interface | 874](#)  
[Check for Fast Ethernet or Gigabit Ethernet Interface Error Statistics | 875](#)

## Verify That the Fast Ethernet or Gigabit Ethernet Interface Is Up

### 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:    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |
| VLAN tagged frames | 0       |          |
| Code violations    | 0       |          |

| Filter statistics:        |   |   |
|---------------------------|---|---|
| Input packet count        | 0 |   |
| Input packet rejects      | 0 |   |
| Input DA rejects          | 0 |   |
| Input SA rejects          | 0 |   |
| Output packet count       |   | 0 |
| Output packet pad count   |   | 0 |
| Output packet error count |   | 0 |

CAM destination filters: 0, CAM source filters: 0

Autonegotiation information:

Negotiation status: Complete

Link partner:

Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:

OK

Local resolution:

Flow control: Symmetric, Remote fault: Link OK

Packet Forwarding Engine configuration:

Destination slot: 4

CoS information:

Direction : Output

| CoS transmit queue |    | Bandwidth |    | Buffer | Priority |
|--------------------|----|-----------|----|--------|----------|
| Limit              | %  | bps       | %  | usec   |          |
| 0 best-effort      | 95 | 950000000 | 95 | 0      | low      |
| none               |    |           |    |        |          |
| 3 network-control  | 5  | 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 112 on page 868](#) presents problem situations and actions for a physical link that is down.

Table 112: 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:
    Receive      Transmit
    Total octets    130624    130624
    Total packets   1281      1281
    Unicast packets  1280      1280
    Broadcast packets      1      1
    Multicast packets      0      0
    CRC/Align errors      0      0
    FIFO errors           0      0
    MAC control frames     0      0
    MAC pause frames       0      0
    Oversized frames       0
    Jabber frames          0
    Fragment frames        0
    VLAN tagged frames     0
    Code violations        0
Filter statistics:
    Input packet count      1281
    Input packet rejects     0
    Input DA rejects        0
    Input SA rejects        0
    Output packet count           1281
    Output packet pad count      0
    Output packet error count    0
    CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
    Local resolution:
        Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
    Destination slot: 4
CoS information:
    Direction : Output

```

| CoS transmit queue<br>Limit |    | Bandwidth |    | Buffer Priority |     |
|-----------------------------|----|-----------|----|-----------------|-----|
|                             | %  | 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

### 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

### Purpose

Persistent interface error statistics indicate that you need to open a case with the Juniper Networks Technical Assistance Center (JTAC).

### Action

To check the local interface for error statistics, use the following Junos OS CLI operational mode command:

```
user@host> show interfaces (fe-fpc/pic/port | ge-fpc/pic/port) extensive
```

### Sample Output

```
user@router> show interfaces ge-4/0/6 extensive
```

```
Physical interface: ge-4/0/6, Enabled, Physical link is Up   Interface index: 144, SNMP ifIndex:
516, Generation: 147
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Enabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags      : Present Running Loop-Detected
  Interface flags:  SNMP-Traps Internal: 0x4000
  Link flags        : None
  CoS queues        : 8 supported, 4 maximum usable queues
  Schedulers        : 0
  Hold-times        : Up 0 ms, Down 0 ms
  Current address:  00:1f:12:fe:c5:2e, Hardware address: 00:1f:12:fe:c5:2e
  Last flapped      : 2015-01-20 23:40:04 PST (00:02:12 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes      :                                0          0 bps
    Output bytes     :                                0          0 bps
    Input packets    :                                0          0 pps
    Output packets   :                                0          0 pps
```

```

IPv6 transit statistics:
  Input  bytes  :                0
  Output bytes  :                0
  Input  packets:                0
  Output packets:                0
Dropped traffic statistics due to STP State:
  Input  bytes  :                0
  Output bytes  :                0
  Input  packets:                0
  Output packets:                0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors:
0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort      0                0                0
  1 expedited-fo     0                0                0
  2 assured-forw     0                0                0
  3 network-cont     0                0                0
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control

Active alarms : None
Active defects : None
MAC statistics:

```

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |

```

VLAN tagged frames          0
Code violations              0
Filter statistics:
  Input packet count        0
  Input packet rejects      0
  Input DA rejects          0
  Input SA rejects          0
  Output packet count       0
  Output packet pad count   0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 4
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
    0 best-effort          95      950000000    95      0      low
none
    3 network-control      5       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

### Purpose

When you suspect a circuit problem, it is important to work with the transport-layer engineer to resolve the problem. The transport-layer engineer may create a loop to the router from various points in the



network. You can then perform tests to verify the connection from the router to that loopback in the network.

### Action

After the transport-layer engineer has created the loop to the router from the network, you must verify the connection from the router to the loopback in the network. Follow Step 2 through Step 8 in [“Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface” on page 860](#). 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 | 878](#)
- [Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface | 879](#)
- [Create a Loopback | 880](#)
- [Verify That the Gigabit Ethernet Interface Is Up | 882](#)
- [Configure a Static Address Resolution Protocol Table Entry | 887](#)
- [Clear Gigabit Ethernet Interface Statistics | 889](#)
- [Ping the Gigabit Ethernet Interface | 889](#)
- [Check for Gigabit Ethernet Interface Error Statistics | 890](#)
- [Diagnose a Suspected Circuit Problem | 894](#)

### Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces

#### Purpose

To use loopback testing to isolate 10, 40, and 100 Gigabit Ethernet interface problems.

#### Action

[Table 111 on page 859](#) provides links and commands for using loopback testing for 10, 40, and 100 Gigabit Ethernet interfaces.

Table 113: Checklist for Using Loopback Testing for 10, 40, and 100 Gigabit Ethernet Interfaces

| Tasks                                                                                        | Command or Action                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>“Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface” on page 879</b> |                                                                                                                                                                                                                                                                                         |
| 1. <a href="#">Create a Loopback on page 861</a>                                             |                                                                                                                                                                                                                                                                                         |
| a. <a href="#">Create a Physical Loopback for a Fiber-Optic Interface on page 861</a>        | Connect the transmit port to the receive port.                                                                                                                                                                                                                                          |
| b. <a href="#">Configure a Local Loopback on page 863</a>                                    | <code>[edit interfaces <i>interface-name</i><br/>(<i>gigether-options</i>)]<br/>set loopback<br/>show<br/>commit</code>                                                                                                                                                                 |
| 2. <a href="#">Verify That the Gigabit Ethernet Interface Is Up on page 882</a>              | <code>show interfaces (xe-fpc/pic/port   et-fpc/pic/port)</code>                                                                                                                                                                                                                        |
| 3. <a href="#">Configure a Static Address Resolution Protocol Table Entry on page 868</a>    | <code>show interfaces (xe-fpc/pic/port  <br/>et-fpc/pic/port)[edit interfaces <i>interface-name</i><br/>unit <i>logical-unit-number</i> family inet address<br/><i>address</i>]<br/>set arp <i>ip-address</i> mac <i>mac-address</i> show<br/>commit<br/>run show arp no-resolve</code> |
| 4. <a href="#">Clear Gigabit Ethernet Interface Statistics on page 889</a>                   | <code>clear interfaces statistics xe-fpc/pic/port  <br/>et-fpc/pic/port</code>                                                                                                                                                                                                          |
| 5. <a href="#">Ping the Gigabit Ethernet Interface on page 889</a>                           | <code>ping <i>remote-IP-address</i> bypass-routing interface<br/>(xe-fpc/pic/port   et-fpc/pic/port) count 100</code>                                                                                                                                                                   |
| 6. <a href="#">Check for Gigabit Ethernet Interface Error Statistics on page 890</a>         | <code>show interfaces (xe-fpc/pic/port   et-fpc/pic/port<br/>) extensive</code>                                                                                                                                                                                                         |
| <b>“Diagnose a Suspected Circuit Problem” on page 877</b>                                    | Perform Steps 2 through 8 from <a href="#">“Diagnose a Suspected Hardware Problem with a Fast Ethernet or Gigabit Ethernet Interface” on page 860</a> .                                                                                                                                 |

## Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface

### Problem

**Description:** When you suspect a hardware problem, take the following steps to help verify if there is a problem.

## Solution

To diagnose a suspected hardware problem with the Ethernet interface, follow these steps:

- [Create a Loopback on page 861](#)
- [Verify That the Gigabit Ethernet Interface Is Up on page 882](#)
- [Configure a Static Address Resolution Protocol Table Entry on page 868](#)
- [Clear Gigabit Ethernet Interface Statistics on page 889](#)
- [Check for Gigabit Ethernet Interface Error Statistics on page 890](#)

## Create a Loopback

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).

1. [Create a Physical Loopback for a Fiber-Optic Interface | 880](#)
2. [Configure a Local Loopback | 881](#)

### *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

[Configure a Local Loopback | 863](#)

[Verify That the Gigabit Ethernet Interface Is Up | 882](#)

[Configure a Static Address Resolution Protocol Table Entry | 868](#)

### **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 gether-options
```

2. Configure the local loopback:

```
[edit interfaces interface-name gether-options
user@router# set loopback
```

3. Verify the configuration:

```
user@router# show
```

For example:

```
[edit interfaces xe-2/0/0 gether-options]
user@router# show
loopback;
```

4. Commit the change:

```
user@router# commit
```

For example:

```
[edit interfaces xe-2/0/0 gether-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 | 882](#)

[Configure a Static Address Resolution Protocol Table Entry | 868](#)

#### SEE ALSO

[Verify That the Gigabit Ethernet Interface Is Up | 882](#)

[Configure a Static Address Resolution Protocol Table Entry | 868](#)

[Clear Gigabit Ethernet Interface Statistics | 889](#)

[Ping the Gigabit Ethernet Interface | 889](#)

[Check for Gigabit Ethernet Interface Error Statistics | 890](#)

## Verify That the Gigabit Ethernet Interface Is Up

### 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
```



```

    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 Degradate :
  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
    Policar: Input: __default_arp_policar__

```



## 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
    Bit errors                Seconds
    Errored blocks            6
  Ethernet FEC Mode : NONE
  Ethernet FEC statistics
    FEC Corrected Errors      Errors
    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 112 on page 868](#) presents problem situations and actions for a physical link that is down.

Table 114: 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:

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

### 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

#### 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
Link Degradate :
  Link Monitoring      : Disable
Interface transmit statistics: Disabled

Logical interface xe-2/0/0.0 (Index 353) (SNMP ifIndex 599) (Generation 175)
Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:
  Input  bytes :      537600
  Output bytes :      539000
  Input  packets:      6400
  Output packets:      6400
Local statistics:
  Input  bytes :      0
  Output bytes :      9800
  Input  packets:      0
  Output packets:      100
Transit statistics:
  Input  bytes :      537600      0 bps
  Output bytes :      529200      0 bps
  Input  packets:      6400      0 pps
  Output packets:      6300      0 pps
Protocol inet, MTU: 1500

```



```

Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 206, Route table: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.108.120.0/30, Local: 10.108.120.1, Broadcast: 10.108.120.3,
Generation: 146
Protocol multiservice, MTU: Unlimited, Generation: 207, Route table: 0
Policer: Input: __default_arp_policer__

```

### Meaning

Check for any error statistics. There should not be any input or output errors. If there are any persistent input or output errors, open a case with the Juniper Networks Technical Assistance Center (JTAC) at [support@juniper.net](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

### Purpose

When you suspect a circuit problem, it is important to work with the transport-layer engineer to resolve the problem. The transport-layer engineer may create a loop to the router from various points in the network. You can then perform tests to verify the connection from the router to that loopback in the network.

### Action

After the transport-layer engineer has created the loop to the router from the network, you must verify the connection from the router to the loopback in the network. Follow Step 2 through Step 8 in [“Diagnose a Suspected Hardware Problem with a Gigabit Ethernet Interface” on page 879](#). 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 | 895](#)
- [Configuring BERT Testing | 897](#)
- [Starting and Stopping a BERT Test | 901](#)

### Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, NxDSO, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own Data Link Layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own Data Link Layer packets, you can assume that the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

The following types of loopback testing are supported by Junos OS:

- DCE local—Loops packets back on the local data circuit-terminating equipment (DCE).
- DCE remote—Loops packets back on the remote DCE.
- Local—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces interface-name** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.

- **Payload**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- **Remote**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 115 on page 896 shows the loopback modes supported on the various interface types.

**Table 115: Loopback Modes by Interface Type**

| Interface                                            | Loopback Modes                           | Usage Guidelines                                                                                                                                                                                                                                         |
|------------------------------------------------------|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet | Local                                    | <i>Configuring Ethernet Loopback Capability</i>                                                                                                                                                                                                          |
| Circuit Emulation E1                                 | Local and remote                         | <i>Configuring E1 Loopback Capability</i>                                                                                                                                                                                                                |
| Circuit Emulation T1                                 | Local and remote                         | <i>Configuring T1 Loopback Capability</i>                                                                                                                                                                                                                |
| E1 and E3                                            | Local and remote                         | <i>Configuring E1 Loopback Capability and Configuring E3 Loopback Capability</i>                                                                                                                                                                         |
| NxDS0                                                | Payload                                  | <i>Configuring NxDS0 IQ and IQE Interfaces, Configuring T1 and NxDS0 Interfaces, Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode), Configuring Fractional E1 IQ and IQE Interfaces, and Configuring Channelized T3 IQ Interfaces</i> |
| Serial (V.35 and X.21)                               | Local and remote                         | <i>Configuring Serial Loopback Capability</i>                                                                                                                                                                                                            |
| Serial (EIA-530)                                     | DCE local, DCE remote, local, and remote | <i>Configuring Serial Loopback Capability</i>                                                                                                                                                                                                            |
| SONET/SDH                                            | Local and remote                         | <i>Configuring SONET/SDH Loopback Capability to Identify a Problem as Internal or External</i>                                                                                                                                                           |

Table 115: Loopback Modes by Interface Type (*continued*)

| Interface | Loopback Modes             | Usage Guidelines                                                                                                                                    |
|-----------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| T1 and T3 | Local, payload, and remote | <i>Configuring T1 Loopback Capability and Configuring T3 Loopback Capability</i><br><br><i>See also Configuring the T1 Remote Loopback Response</i> |

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 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-2e7           Pattern is 2^7 - 1
pseudo-2e9-o153      Pattern is 2^9 - 1 (per O.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 O.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 0.153 standard)    |
| 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-2e20-o153              | Pattern is $2^{20} - 1$ (per 0.153 standard) |
| pseudo-2e23-o151              | Pattern is $2^{23}$ (per 0.151 standard)     |

Table 116 on page 900 shows the BERT capabilities for various interface types.

**Table 116: BERT Capabilities by Interface Type**

| Interface                                     | T1 BERT          | T3 BERT            | Comments                                                                                                                     |
|-----------------------------------------------|------------------|--------------------|------------------------------------------------------------------------------------------------------------------------------|
| 12-port T1/E1 Circuit Emulation               | Yes (ports 0-11) | —                  | <ul style="list-style-type: none"> <li>Limited algorithms</li> </ul>                                                         |
| 4-port Channelized OC3/STM1 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 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> |

Table 116: BERT Capabilities by Interface Type (*continued*)

| Interface                          | T1 BERT            | T3 BERT                     | Comments                                                                                                                                                                              |
|------------------------------------|--------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Channelized 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 and Multichannel 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 \(Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port\) | 1580](#)

## 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 | 902](#)
- [Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm | 903](#)
- [Fast Ethernet and Gigabit Ethernet Counters | 905](#)

### Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

#### Purpose

To locate LINK alarm and major counters associated with Fast Ethernet and Gigabit Ethernet interfaces.

#### Action

[Table 117 on page 903](#) provides links and commands for locating LINK alarm and major counters for Fast Ethernet and Gigabit Ethernet interfaces.

Table 117: Checklist for Locating Fast Ethernet and Gigabit Ethernet Alarms and Counters

| Tasks                                                                                            | Command or Action                                                          |
|--------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| <a href="#">“Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm” on page 903</a> | <code>show interfaces (fe-fpc/pic/port   ge-fpc/pic/port) extensive</code> |
| <a href="#">“Fast Ethernet and Gigabit Ethernet Counters” on page 905</a>                        |                                                                            |

## SEE ALSO

*Ethernet Interfaces User Guide for Routing Devices*

**Display the Fast Ethernet or Gigabit Ethernet Interface LINK Alarm****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” on page 859](#) 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 | 902](#)

[Fast Ethernet and Gigabit Ethernet Counters | 905](#)

*Ethernet Interfaces User Guide for Routing Devices*

## Fast Ethernet and Gigabit Ethernet Counters

### Problem

**Description:** [Table 118 on page 906](#) 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 118: 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 aborts 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.                                                                          |
| <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 aborts and FCS errors.                                                                                       |                                                                                                                                                                                                                              |

Table 118: 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). |
| <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.                                                                                                                                                                                                     |

Table 118: Major Fast Ethernet and Gigabit Ethernet Counters (continued)

| Counter                                            | Description                                                                                                            | Reason for 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 | 902](#)

*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

### 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 909](#) section for recovery options.

1.

**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 | 3
- 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 | 910
- Guidelines to perform Bidirectional Diagnostics using Remote Loopback | 914
- Clearing the Interface Statistics | 917

## 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) | -       |

| MIC Type        | PRBS Test Engine                     | Remarks                                                                                                                                                                                                 |
|-----------------|--------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |
| 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

|                                                                   |
|-------------------------------------------------------------------|
| <a href="#">prbs-test-start</a>   <a href="#">1473</a>            |
| <a href="#">prbs-test-stop</a>   <a href="#">1475</a>             |
| <a href="#">show interfaces prbs-stats</a>   <a href="#">1842</a> |
| <a href="#">clear interfaces statistics</a>                       |

## Guidelines to perform Bidirectional Diagnostics using Remote Loopback

Configure loopback on the remote end of the link. Test pattern generation and analysis is performed on the local end.

The PRBS statistics report on the receiving end reflects link quality.

Following are the steps to collect and view the PRBS statistics:

1. Enable remote loopback.

```
user@host2> set interfaces et-1/1/4 gigether-options loopback-remote
```

**NOTE:** You must disable Forward Error Correction (FEC), if loopback is configured in the router with JNP-MIC1 MIC at the remote end.

```
user@host2> show interfaces et-1/1/4 | display set
```

```
set interfaces et-1/1/4 gigether-options fec none
```

2. Start a TX (direction 0) by issuing the following command:

```
[edit]
```

```
user@host1> test interface et-0/1/2 prbs-test-start pattern-type 31 direction 0 flip 0
```

After executing the command, you can check the link status by executing `show interfaces terse et-0/1/*`:

```
Interface Admin    Link Proto Local Remote
et-0/1/2 up        down
```

3. Start a RX (direction 1) by issuing the following command on the same host.

```
[edit]
```

```
user@host1> test interface et-0/1/2 prbs-test-start pattern-type 31 direction 1 flip 0
```

**NOTE:** There is a change in the direction as 1.

4. After starting the statistics collection, you can view the collected statistics at RX by issuing the following command:

```
[edit]
user@host1> show interfaces interface-name prbs-stats
```

For example:

Checking PRBS statistics at RX:

```
user@host1> show interfaces et-0/1/2 prbs-stats
```

```
PRBS Statistics : Enabled

Lane 0 : State : Pass, Error count : 0
Lane 1 : State : Pass, Error count : 0
Lane 2 : State : Pass, Error count : 0
Lane 3 : State : Pass, Error count : 0
```

The PRBS test is successful, if the state is pass with error count 0.

5. Stop the PRBS statistics collection by issuing the following command:

```
user@host1> test interface interface-name prbs-test-stop direction 1
user@host1> test interface interface-name prbs-test-stop direction 0
```

For example:

To stop the PRBS at RX:

```
user@host1> test interface et-0/1/2 prbs-test-stop direction 1
```

Verify the statistics at TX by executing the following command:

```
show interfaces et-0/1/2 prbs-stats
```

```
PRBS Statistics : Disabled
```

To stop the PRBS at TX:

```
user@host1> test interface et-0/1/2 prbs-test-stop direction 0
```

After executing the command, you can check the link status at RX by executing `show interfaces terse et-0/1/2`:

```
Interface Admin    Link Proto Local Remote
et-0/1/2 up       up
```

This command only disables the statistics collection and does not clear the statistics collected. To clear the collected statistics, issue the clear interfaces statistics command.

```
user@host1> clear interfaces statistics et-0/1/2
```

### Interface Card Specific differences

While collecting statistics, the JNP-MIC1-MACSEC and JNP-MIC1 MICs behave differently:

- On JNP-MIC1-MACSEC MIC, if RX is not latching to any PRBS signal, then the "state" in the **show interfaces interface-name prbs-stats** displays as "Disabled" with Error count as 0, where the JNP-MIC1 MIC displays as failed with MAX error count.

For Example: On JNP-MIC1-MACSEC MIC

```
user@host> test interface et-0/1/10 prbs-test-start pattern-type 31 direction 1 flip 0
user@host> show interfaces et-0/1/10 prbs-stats
```

```
PRBS Statistics : Enabled
Lane 0 : State : Disabled, Error count : 0
Lane 1 : State : Disabled, Error count : 0
Lane 2 : State : Disabled, Error count : 0
Lane 3 : State : Disabled, Error count : 0
```

For Example: On JNP-MIC1 MIC

```
user@host> test interface et-0/0/1 prbs-test-start pattern-type 31 direction 1 flip 0
user@host> show interfaces et-0/0/1 prbs-stats
```

```
PRBS Statistics : Enabled
Lane 0 : State : Fail, Error count : 4294967295
Lane 1 : State : Fail, Error count : 4294967295
Lane 2 : State : Fail, Error count : 4294967295
Lane 3 : State : Fail, Error count : 4294967295
```

- If any mismatch is encountered between the pattern-type and flip between TX and RX, the maximum error counts are observed in case of JNP-MIC1 MIC and state "disabled" in case of JNP-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 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 together-options loopback-remote
set interfaces et-1/1/1 together-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 | 1473](#)

[prbs-test-stop | 1475](#)

[show interfaces prbs-stats | 1842](#)



# 5

PART

## Configuration Statements and Operational Commands

---

Configuration Statements (OTN) | **923**

Configuration Statements (OAM-CFM) | **989**

Configuration Statements | **1081**

Operational Commands | **1427**

---



## Configuration Statements (OTN)

### IN THIS CHAPTER

- alarm (optics-options) | 925
- backward-frr-enable | 926
- ber-threshold-clear | 928
- ber-threshold-signal-degrade | 931
- bypass | 934
- bytes (otn-options) | 935
- cfp-to-et | 936
- fec | 937
- fec (gigether) | 940
- fixed-stuff-bytes | 941
- high-polarization | 942
- interval | 943
- is-ma | 945
- laser-enable | 946
- line-loopback | 947
- local-loopback | 948
- monitor-end-point | 949
- no-odu-backward-frr-enable | 950
- no-odu-signal-degrade-monitor-enable | 951
- number-of-frames | 952
- oc192 | 953
- odu-delay-management | 954
- odu-backward-frr-enable | 955
- odu-signal-degrade | 956
- odu-signal-degrade-monitor-enable | 958
- odu-ttim-action-enable | 959
- otu-ttim-action-enable | 960
- otu4 | 961

- [pass-through](#) | 962
- [prbs](#) | 963
- [preemptive-fast-reroute](#) | 964
- [rate](#) | 965
- [remote-loop-enable](#) | 966
- [signal-degrade](#) | 968
- [signal-degrade-monitor-enable](#) | 969
- [start-measurement](#) | 971
- [tca](#) | 972
- [transport-monitoring](#) | 975
- [trigger](#) | 976
- [tti](#) | 981
- [tx-power](#) | 982
- [warning](#) | 983
- [wavelength](#) | 984

## alarm (optics-options)

### Syntax

```
alarm low-light-alarm {  
    (link-down | syslog);  
}
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

### Release Information

Statement introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### Description

Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold.

### 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.

## RELATED DOCUMENTATION

[Configuring Link Down Notification for Optics Options Alarm or Warning | 237](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

## backward-frr-enable

### Syntax

```
(backward-frr-enable | no-backward-frr-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### 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 498](#) 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.

#### RELATED DOCUMENTATION

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 498](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## ber-threshold-clear

### Syntax

```
ber-threshold-clear value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

### 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 498](#) 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 119 on page 929 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 119: 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  $x\text{E}-n$ , 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 119 on page 929](#) 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 498](#).

**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  $1.0\text{E}-n$  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.

### RELATED DOCUMENTATION

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 498](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## ber-threshold-signal-degrade

### 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]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

### 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 498](#) 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 120 on page 932 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 120: 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                                         |
| 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  $x\text{E}-n$ , where  $x$  is the mantissa and  $n$  is the exponent. For example,  $4.5\text{E}-3$ .

**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 120 on page 932](#).

**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  $1.0\text{E}-n$  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 498](#).

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 498](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## bypass

### Syntax

```
(bypass | no-bypass);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### 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

**bypass**—Pass the DM value through a node.

**no-bypass**—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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)



## bytes (otn-options)

### Syntax

```
bytes transmit-payload-type value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## cfp-to-et

### Syntax

```
cfp-to-et;
```

### Hierarchy Level

```
[edit chassis fpc slot]
```

### Release Information

Statement introduced in Junos OS Release 19.2R1-S1.

### 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.

### RELATED DOCUMENTATION

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

## fec

### List of Syntax

[Syntax \(M Series, MX Series, PTX Series\) on page 937](#)

[Syntax \(ACX6360\) on page 937](#)

### 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]
```

### 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.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 18.3R1 for ACX6360 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

## RELATED DOCUMENTATION

---

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

---

[Supported Forward Error Correction Modes on MX Series Routers | 502](#)

---

[Supported Forward Error Correction Modes on PTX Series Routers | 503](#)

---

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 498](#)

---

[Supported Forward Error Correction Modes on ACX6360 Router | 504](#)

## fec (gigether)

### Syntax

```
fec (fec91 |fec74 |none)
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

### Release Information

Statement introduced in Junos OS Release 17.3R1.

### Description

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. Delete the statement and commit the configuration to return to the default behavior.

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.

To disable the FEC mode, use the command **set interfaces *interface-name* gigether-options fec none**.

To re-enable the FEC mode, use the commands **set interfaces *interface-name* gigether-options fec fec91**, or **delete interfaces *interface-name* gigether-options fec none**.

To check FEC status, use the command **show interfaces *interface-name***.

### 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.

## fixed-stuff-bytes

### Syntax

```
(fixed-stuff-bytes | no-fixed-stuff-bytes);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

### Description

Enable or disable fixed stuff bytes.

### Default

By default, no fixed stuff bytes are set.

### Options

**fixed-stuff-bytes**—Fixed stuff bytes 11.0957 Gbps.

**no-fixed-stuff-bytes**—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.

## RELATED DOCUMENTATION

---

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## high-polarization

### Syntax

```
high-polarization;
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

### Release Information

Statement introduced in Junos OS Release 18.2R1 for ACX6360 routers.

### 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.

## RELATED DOCUMENTATION

[Supported Optics Options on ACX6360 and ACX5448-D Routers](#) | 489



## interval

### Syntax

```
interval value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options signal-degrade]
[edit interfaces interface-name otn-options odu-signal-degrade]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

### 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 498](#) 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## is-ma

### Syntax

```
(is-ma | no-is-ma);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### Description

Specify whether masked alarms are enabled or disabled.

### Default

If you omit the is-ma statement, masked alarms are disabled.

### Options

**is-ma**—Enable masked alarms.

**no-is-ma**—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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## laser-enable

### Syntax

```
(laser-enable | no-laser-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### Description

Specify whether lasers are enabled or disabled.

### Default

If you omit the laser-enable statement, lasers are disabled.

### Options

**laser-enable**—Enable lasers.

**no-laser-enable**—Do not enable lasers.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## line-loopback

### Syntax

```
(line-loopback-enable | no-line-loopback);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### Description

Specify whether line-loopback is enabled or disabled.

### Default

If you omit the line-loopback-enable statement, line-loopback is disabled.

### Options

**line-loopback-enable**—Enable line-loopback.

**no-line-loopback**—Disable line-loopback.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## local-loopback

### Syntax

```
(local-loopback-enable | no-local-loopback);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### Description

Specify whether local-loopback is enabled or disabled.

### Default

If you omit the local-loopback-enable statement, local-loopback is disabled.

### Options

**local-loopback-enable**—Enable local-loopback.

**no-local-loopback**—Disable local-loopback.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 424

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 512

## monitor-end-point

### Syntax

```
(monitor-end-point | no-monitor-end-point);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### Description

Originate or do not originate the connection monitor end point.

### Default

By default, do not originate the connection monitor end point.

### Options

**monitor-end-point**—Originate the connection monitor end point.

**no-monitor-end-point**—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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## no-odu-backward-frr-enable

### Syntax

```
no-odu-backward-frr-enable;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Release Information

Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

### 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)



## no-odu-signal-degrade-monitor-enable

### Syntax

```
no-odu-signal-degrade-monitor-enable;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Release Information

Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

### 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.

### RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## number-of-frames

### Syntax

```
number-of-frames value;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 424

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 512

## oc192

### Syntax

```
oc192;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

### Release Information

Statement introduced in Junos OS Release 13.3 for MX Series routers.

### 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.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## odu-delay-management

### 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]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## odu-backward-frr-enable

### Syntax

```
odu-backward-frr-enable;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Release Information

Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

### 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## odu-signal-degrade

### Syntax

```
odu-signal-degrade {
  ber-threshold-clear;
  ber-threshold-signal-degrade;
  interval
}
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

**RELATED DOCUMENTATION**

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 424

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 512

## odu-signal-degrade-monitor-enable

### Syntax

```
odu-signal-degrade-monitor-enable;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Release Information

Statement introduced in Junos OS Release 14.1R2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 14.2 for P2-100GE-OTN PIC in PTX5000 routers.

Statement introduced in Junos OS Release 15.1F6 for MPC5E, MIC6-100G-CFP2 on MPC6E on MX240, MX480, MX960, MX2010, and MX2020 routers.

### 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)



## odu-ttim-action-enable

### Syntax

```
(odu-ttim-action-enable | no-odu-ttim-action-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### 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

**odu-ttim-action-enable**—Enable consequent action for ODU TTIM.

**no-odu-ttim-action-enable**—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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## otu-ttim-action-enable

### Syntax

```
(otu-ttim-action-enable | no-otu-ttim-action-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### 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

**otu-ttim-action-enable**—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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## otu4

### Syntax

```
otu4;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 13.3 for MX Series routers.

### 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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 424

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 512

## pass-through

### Syntax

```
(pass-through | no-pass-through);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options rate]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

### Description

Enable or disable OTN pass-through mode.

### Default

By default, OTN pass-through mode is disabled.

### Options

**no-pass-through**—Do not enable OTN pass-through mode.

**pass-through**—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.

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## prbs

### Syntax

```
(prbs | no-prbs);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### Description

Specify whether OTN payload Pseudo-Random Binary Sequence (PBRs) is enabled or disabled.

### Default

By default, OTN payload prbs is disabled.

### Options

**prbs**—Enable OTN payload PBRs.

**no-prbs**—Disable OTN payload PBRs.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 424

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 512

## preemptive-fast-reroute

### 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]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 13.3 for MX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## rate

### Syntax

```
rate {
  (fixed-stuff-bytes | no-fixed-stuff-bytes);
  otu4; oc192;
  (pass-through | no-pass-through);
}
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### 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.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

### 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.

### RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## remote-loop-enable

### Syntax

```
(remote-loop-enable | no-remote-loop-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Release Information

Statement introduced in Junos OS Release 17.1 for 100-Gigabit OTN DWDM PIC with CFP2 on PTX3000 and PTX5000 routers.

Statement introduced in Junos OS Release 17.1 for 100-Gigabit OTN DWDM MIC with CFP2-ACO on MX240, MX480, MX960, MX2010, and MX2020 routers with MPC3E and MPC3E-NG.

### 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.

## RELATED DOCUMENTATION



[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[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 | 521](#)

---

[Configuring OTN Interfaces on PTX-5-100G-WDM PIC | 526](#)

---

[Disabling ODU Path Delay Measurement | 541](#)

---

[Enabling ODU Path Delay Measurement | 539](#)

---

[Understanding ODU Path Delay Measurement | 537](#)

---

[Understanding the MIC3-100G-DWDM MIC | 437](#)

---

[Understanding the PTX-5-100G-WDM PIC | 441](#)

---

## signal-degrade

### Syntax

```
signal-degrade {
  ber-threshold-clear value;
  ber-threshold-signal-degrade value;
  interval value;
}
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 13.3 for MX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## signal-degrade-monitor-enable

### Syntax

```
(signal-degrade-monitor-enable | no-signal-degrade-monitor-enable);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options preemptive-fast-reroute]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### 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 498](#) 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.

**RELATED DOCUMENTATION**

---

[Understanding Pre-FEC BER Monitoring and BER Thresholds | 498](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## start-measurement

### Syntax

```
(no-start-measurement | start-measurement);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options odu-delay-management]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### Description

Start or do not start a delay measurement (DM) session.

### Default

By default, do not start a DM session.

### Options

**no-start-measurement**—Do not start a DM session.

**start-measurement**—Start a DM session.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## tca

### 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]
```

### Release Information

Statement introduced in Junos OS Release 14.2 on the PTX Series.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.



## RELATED DOCUMENTATION

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

[optics-options | 1297](#)

## transport-monitoring

### Syntax

```
transport-monitoring;
```

### Hierarchy Level

```
[edit interfaces]
```

### Release Information

Statement introduced in Junos OS Release 14.2 for PTX5000 and PTX3000 routers.

### 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.

## RELATED DOCUMENTATION

[show interfaces transport pm | 1844](#)

## trigger

### Syntax

```
trigger trigger-identifier (hold-time hold-time-value | ignore);
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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-loss**—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.

#### RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## tti

### Syntax

```
tti tti-identifier;
```

### Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 13.3 for MX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

### RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM | 512](#)

## tx-power

### Syntax

```
tx-power dbm;
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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.

## RELATED DOCUMENTATION

[Ethernet DWDM Interface Wavelength Overview | 505](#)

[optics-options | 1297](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)



## warning

### Syntax

```
warning low-light-warning {
    (link-down | syslog);
}
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

### Release Information

Statement introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### Description

Specifies the action to take if the receiving optics signal is below the optics low-light warning threshold.

### 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.

## RELATED DOCUMENTATION

[Configuring Link Down Notification for Optics Options Alarm or Warning | 237](#)

[optics-options | 1297](#)

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

## wavelength

### Syntax

```
wavelength nm;
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series routers.

### 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.

**RELATED DOCUMENTATION**

---

[Ethernet DWDM Interface Wavelength Overview | 505](#)

---

[Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength | 505](#)

---

[show interfaces diagnostics optics \(Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port\) | 1580](#)

## Configuration Statements (OAM-CFM)

### IN THIS CHAPTER

- 802.3ad | 992
- action-profile (Applying to CFM) | 993
- action-profile (Defining for CFM) | 994
- action-profile (MEP) | 996
- ais-trigger-condition | 997
- all-defects | 998
- auto-discovery | 999
- avg-fd-twoway-threshold | 1000
- avg-ifdv-twoway-threshold | 1002
- avg-flr-forward-threshold | 1004
- avg-flr-backward-threshold | 1006
- calculation-weight | 1008
- clear-action (CFM) | 1009
- connection-protection-tlv | 1010
- continuity-check | 1011
- convey-loss-threshold | 1012
- cross-connect-ccm | 1013
- cycle-time | 1014
- data-tlv-size | 1015
- default-actions | 1016
- delay | 1017
- delegate-server-processing | 1018
- delay-variation | 1019
- detect-loc | 1020
- direction | 1022
- enhanced-cfm-mode | 1023
- erroneous-ccm | 1024
- event (CFM) | 1025

- flap-trap-monitor | **1026**
- hardware-assisted-timestamping | **1028**
- hardware-assisted-keepalives | **1029**
- hold-interval (OAM) | **1031**
- instance | **1032**
- interface (IEEE 802.1ag OAM Connectivity-Fault Management) | **1033**
- interface (OAM Link-Fault Management) | **1034**
- interface-down | **1035**
- interface-status-tlv | **1036**
- interface-status-send-rdi | **1037**
- interval | **1038**
- interval (CFM MEP) | **1040**
- iteration-count | **1041**
- iteration-period | **1042**
- level | **1043**
- level (CFM MEP) | **1044**
- linktrace | **1045**
- log-and-generate-ais | **1046**
- loss-threshold | **1047**
- lowest-priority-defect | **1048**
- maintenance-association | **1049**
- maintenance-domain | **1051**
- measurement-interval | **1053**
- measurement-type | **1054**
- mep | **1055**
- mip-half-function | **1057**
- name-format | **1058**
- no-aggregate-delegate-processing | **1059**
- path-database-size | **1060**
- performance-monitoring | **1061**
- policer (CFM Global) | **1063**
- policer (CFM Session) | **1064**
- port-status-tlv | **1065**
- priority (Protocols OAM) | **1066**



- [priority \(CFM MEP\) | 1067](#)
- [priority \(OAM Connectivity-Fault Management\) | 1068](#)
- [protocol \(Server MEP\) | 1069](#)
- [protect-maintenance-association \(OAM\) | 1070](#)
- [rdi | 1071](#)
- [receive-ais | 1072](#)
- [remote-maintenance-association \(OAM\) | 1073](#)
- [remote-mep | 1074](#)
- [sendid-tlv | 1075](#)
- [short-name-format | 1076](#)
- [sla-iterator-profile | 1077](#)
- [sla-iterator-profiles | 1079](#)

## 802.3ad

### Syntax

```
802.3ad {
    primary | backup;
    ae interface-number ;
    lacp {
        port-priority priority-number;
    }
    link-index index-number
    distribution-list distribution-list-number
}
```

### Hierarchy Level

```
[edit interfaces interface-name fastether-options],
[edit interfaces interface-name gigether-options]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

**primary** and **backup** options added in Junos OS Release 8.3.

### Description

Specify aggregated Ethernet logical interface number.

### Options

**bundle**—Join an aggregated Ethernet interface.

**ae *interface-number***—Aggregated Ethernet logical interface number. For MX Series routers running Junos release 14.2R3 and later you can configure a maximum of 1000 aggregated interfaces. On MX2010 and MX2020 routers you can configure a maximum of 800 aggregated interfaces.

**primary | backup**—For link protection configurations, specify the link as primary link or backup link for egress traffic.

**lacp**—Configure Link Aggregation Control Protocol. Specify the port priority in the range 0 through 65535. Default port-priority is 127.

**link-index**—Specify the desired child link index within the aggregated Ethernet Interface. Index number of the logical interface reflects its initialization sequence.

**distribution-list**—For targeted distribution, specify the distribution list to which the interface belongs.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring an Aggregated Ethernet Interface | 64](#)

[Configuring Aggregated Ethernet Link Protection | 133](#)

## action-profile (Applying to CFM)

### Syntax

```
action-profile profile-name;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name mep mep-id remote-mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

Identify the action profile to use.

### Options

***profile-name***—Name of the action profile to use.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages | 564](#)

## action-profile (Defining for CFM)

### Syntax

```

action-profile profile-name {
  event {
    ais-trigger-condition {
      adjacency-loss;
      all-defects;
      cross-connect-ccm;
      erroneous-ccm;
      receive-ais;
    }
    interface-status-tlv (down | lower-layer-down);
    port-status-tlv blocked;
    rdi;
  }
  action {
    interface-down;
    log-and-generate-ais {
      interval(1m | 1s);
      level value;
      priority value;
    }
  }
  default-actions {
    interface-down;
  }
}

```

### Hierarchy Level

[edit protocols oam ethernet [connectivity-fault-management](#)]

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

Configure a name and default action for an action profile.

### Options

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

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

**Required Privilege Level**

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

|  |                                                                                              |
|--|----------------------------------------------------------------------------------------------|
|  | <a href="#">Configuring a CFM Action Profile to Specify CFM Actions for CFM Events   589</a> |
|  | <a href="#">default-actions   1016</a>                                                       |
|  | <a href="#">event (CFM)   1025</a>                                                           |
|  | <a href="#">interface-down   1035</a>                                                        |

## action-profile (MEP)

### Syntax

```
action-profile action-profile-name;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association  
  ma-name mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Attach the configured action profile to the MEP depending on the hierarchy level.

### Options

***action-profile-name***—Name of the action profile that is configured for the CFM MEP and the server MEP.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## ais-trigger-condition

### Syntax

```
ais-trigger-condition {  
    adjacency-loss;  
    all-defects;  
    cross-connect-ccm;  
    erroneous-ccm;  
    receive-ais;  
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the defect conditions that generate an alarm indication signal (AIS).

### Options

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview](#) | 810

[Configuring ETH-AIS on a CFM MEP](#) | 816

## all-defects

### Syntax

```
all-defects;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event  
  ais-trigger-condition]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the defect condition that raises an alarm indication signal when any or all possible defects occur in the maintenance domain level.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)



## auto-discovery

### Syntax

```
auto-discovery;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

Enable the MEP to accept continuity check messages from all remote MEPs.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564

## avg-fd-twoway-threshold

### Syntax

```
avg-fd-twoway-threshold avg-fd-twoway-threshold-value{
  flap-trap-monitor seconds;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
  profile-name]
```

### Release Information

Statement introduced in Junos OS Release 15.1 for MX Series routers.

### Description

Configure the threshold value for average frame delay, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay is exceeded, an SNMP trap is generated for ETH-DM. SNMP traps are triggered if you configure either the average frame-delay threshold or the average inter-frame delay variance threshold. If you do not configure either the frame-delay threshold or the frame delay variation threshold, no SNMP traps are generated. You can configure these threshold values only with a two-way ETH-DM SLA iterator.

Frame delay refers to the difference, in microseconds, between the time a frame is sent and when it is received. Frame delay variation refers to the difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.” For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

### Options

***avg-fd-twoway-threshold-value***—Threshold value for average frame delay, in microseconds, for two-way ETH-DM.

**Range:** 1 through 4294967295 microseconds

***flap-trap-monitor seconds***—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

**Range:** 1 through 360 seconds

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[flap-trap-monitor](#) | 1026

---

[Configuring an Iterator Profile](#) | 782

---

*Configuring an Iterator Profile on a Switch (CLI Procedure)*

## avg-ifdv-twoway-threshold

### Syntax

```
avg-ifdv-twoway-threshold avg-ifdv-twoway-threshold-value; {
    flap-trap-monitor seconds; }
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
    profile-name]
```

### Release Information

Statement introduced in Junos OS Release 15.1 for MX Series routers.

### Description

Configure the threshold value for average frame delay variation, in microseconds, for two-way Ethernet frame delay measurement (ETH-DM). When the configured threshold for average frame delay variation is exceeded, an SNMP trap is generated for ETH-DM. SNMP traps are triggered if you configure either the average frame-delay threshold or the average inter-frame delay variance threshold. If you do not configure either the frame-delay threshold or the frame delay variation threshold, no SNMP traps are generated. You can configure these threshold values only with a two-way ETH-DM SLA iterator.

Frame delay refers to the difference, in microseconds, between the time a frame is sent and when it is received. Frame delay variation refers to the difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.” For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics. In two-way ETH-DM mode, frame delay and frame delay variation values are based on the time difference between when the initiator MEP transmits a request frame and receives a reply frame from the responder MEP, subtracting the time elapsed at the responder MEP.

### Options

***avg-ifdv-twoway-threshold-value***—Threshold value for average frame delay variation, in microseconds, for two-way ETH-DM.

**Range:** 1 through 4294967295 microseconds

***flap-trap-monitor seconds***—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

**Range:** 1 through 360 seconds

### Required Privilege Level

**configure**—To enter configuration mode.

control—To modify any configuration.

#### RELATED DOCUMENTATION

---

[flap-trap-monitor](#) | **1026**

---

[Configuring an Iterator Profile](#) | **782**

---

*Configuring an Iterator Profile on a Switch (CLI Procedure)*

## avg-flr-forward-threshold

### Syntax

```
avg-flr-forward-threshold avg-flr-forward-threshold-value {
  flap-trap-monitor seconds; }
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
  profile-name]
```

### Release Information

Statement introduced in Junos OS Release 15.1 for MX Series routers.

### Description

Configure the threshold value for average frame loss ratio, in milli-percent, in the upstream or forward direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average forward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. SNMP traps are triggered if you configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold. If you do not configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold, no SNMP traps are generated. You can configure these threshold values with an SLA iterator for ETH-SLM and ETH-LM.

ETH-SLM is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic.

### Options

***avg-flr-forward-threshold-value***—Threshold value for average frame loss ratio in the forward or upstream direction, in milli-percent, for ETH-SLM and ETH-LM.

**Range:** 1 through 100000 milli-percent

***flap-trap-monitor seconds***—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

**Range:** 1 through 360 seconds

### Required Privilege Level

configure—To enter configuration mode.

control—To modify any configuration.

## RELATED DOCUMENTATION

[flap-trap-monitor](#) | 1026

---

[Configuring an Iterator Profile](#) | 782

---

*Configuring an Iterator Profile on a Switch (CLI Procedure)*

## avg-flr-backward-threshold

### Syntax

```
avg-flr-backward-threshold avg-flr-backward-threshold-value{
  flap-trap-monitor seconds; }
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles
  profile-name]
```

### Release Information

Statement introduced in Junos OS Release 15.1 for MX Series routers.

### Description

Configure the threshold value for average frame loss ratio, in milli-percent, in the backward or downstream direction for Ethernet loss measurement (ETH-LM) and Ethernet synthetic loss measurement (ETH-SLM). When the configured threshold for average backward frame loss ratio is exceeded, an SNMP trap is generated for ETH-LM and ETH-SLM. SNMP traps are triggered if you configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold. If you do not configure either the average backward frame loss ratio threshold or the average forward frame loss ratio threshold, no SNMP traps are generated. You can configure these threshold values with an SLA iterator for ETH-SLM and ETH-LM.

ETH-SLM is an application that enables the calculation of frame loss by using synthetic frames instead of data traffic. This mechanism can be considered as a statistical sample to approximate the frame loss ratio of data traffic.

### Options

***avg-flr-backward-threshold-value***—Threshold value for average frame loss ratio in the backward or downstream direction, in milli-percent, for ETH-SLM and ETH-LM.

**Range:** 1 through 100000 milli-percent

***flap-trap-monitor seconds***—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPm ThresholdFlapAlarm notification to the network management system (NMS).

**Range:** 1 through 360 seconds

### Required Privilege Level

configure—To enter configuration mode.

control—To modify any configuration.



## RELATED DOCUMENTATION

[flap-trap-monitor](#) | 1026

---

[Configuring an Iterator Profile](#) | 782

---

*Configuring an Iterator Profile on a Switch (CLI Procedure)*

## calculation-weight

### Syntax

```
calculation-weight {  
    delay delay-value;  
    delay-variation delay-variation-value;  
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles  
    profile-name]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

### Description

Configure the calculation weight for delay and delay variation.

**NOTE:** This option is applicable only for two-way delay measurement.

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

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Configuring an Iterator Profile | 782](#)

[Configuring an Iterator Profile on a Switch \(CLI Procedure\)](#)

[delay | 1017](#)

[delay-variation | 1019](#)

## clear-action (CFM)

### Syntax

```
clear-action {  
  interface-down peer-interface;  
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name ]
```

### Release Information

Statement introduced in Junos OS Release 11.4.

### Description

Clear the action or actions to be taken when the connectivity fault management event occurs. You cannot configure multiple actions at this time. Only one action can be configured. This limitation affects both the **action** and **clear-action** statements.

### Options

***peer-interface***—Name of the peer interface.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring a CFM Action Profile to Specify CFM Actions for CFM Events](#) | 589

## connection-protection-tlv

### Syntax

```
connection-protection-tlv;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name continuity-check]
```

### Description

Includes connection protection OUI TLV in continuity check messages (CCM) .The TLV is responsible for carrying the flag information within CCM PDUs. Though this OUI TLV will be included in the CCM frames by provider edge devices, the value is updated by the provider routers in case the traffic to the other end of the network is forwarded by the facility protection tunnel .

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[connectivity-fault-management](#) | [1132](#)

[Configuring MAC Flush Message Processing in CET Mode](#) | [642](#)

[Example: Configuring an Action Profile Based on Connection Protection TLVs](#) | [645](#)

## continuity-check

### Syntax

```
continuity-check {
  convey-loss-threshold;
  hold-interval minutes;
  interface-status-tlv;
  interval (10m | 10s | 1m | 1s| 100ms | 10ms| 3.3ms);
  loss-threshold number;
  port-status-tlv;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

Specify continuity check protocol options.

**NOTE:** When the continuity-check interval is set to 3.3ms, the actual interval is 4ms (250 pps).

### Options

**convey-loss-threshold**—Enable loss-threshold-tlv transmission.

**hold-interval** *minutes*—Specify the continuity check hold-interval, in minutes.

**interface-status-tlv**—Enable interface-status-tlv transmission.

**interval** (*10m | 10s | 1m | 1s| 100ms | 10ms | 3.3ms*)—Specify the continuity check interval.

**loss-threshold** *minutes*—Specify the loss-threshold, in minutes.

**port-status-tlv**—Enable port-status-tlv transmission.

### Required Privilege Level

**interface**—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring Continuity Check Protocol Parameters for Fault Detection](#) | 563

## convey-loss-threshold

### Syntax

```
convey-loss-threshold;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
maintenance-association ma-name convey-loss-threshold]
```

### Description

Enable loss-threshold-tlv transmission.

During a unified ISSU, the control plane may go down for several seconds and cause CFM continuity check packets to get dropped. This may cause the remote maintenance endpoint (MEP) to detect a connectivity loss and mark the MEP as down. To keep the MEP active during a unified ISSU, the loss threshold TLV communicates the minimum threshold value the receiving MEP requires to keep the MEP active. The receiving MEP parses the TLV and updates the loss threshold value, but only if the new threshold value is greater than the locally configured threshold value. You can control the transmission of the loss threshold TLV in continuity check messages PDUs. The **convey-loss-threshold** statement specifies that the loss threshold TLV must be transmitted as part of the continuity check messages. If the statement is not specified, continuity check messages transmit this TLV only when a unified ISSU is in progress.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## cross-connect-ccm

### Syntax

```
cross-connect-ccm;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event  
  ais-trigger-condition]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the defect condition that raises an alarm indication signal when any cross-connect continuity check messages (CCMs) are received by the MEP.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## cycle-time

### Syntax

```
cycle-time cycle-time-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles  
  profile-name]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

### Description

Configure the time (in milliseconds) taken between back-to-back transmissions of SLA frames for a single connection.

### Options

**cycle-time-value**—Cycle time value in milliseconds.

**Range:** 10 through 3,600,000

**Default:** 1000

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Configuring an Iterator Profile | 782](#)

*Configuring an Iterator Profile on a Switch (CLI Procedure)*



## data-tlv-size

### Syntax

```
data-tlv-size size;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name  
maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

### Description

Configure the size of the data TLV portion of the Y.1731 data frame.

### Options

**size**—Size of the data TLV portion of the Y.1731 data frame.

**NOTE:** This option is applicable only for two-way delay measurement.

**Range:** 1 through 1400 bytes

**Default:** 1

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

### RELATED DOCUMENTATION

[sla-iterator-profile](#) | 1077

[Configuring a Remote MEP with an Iterator Profile](#) | 794

## default-actions

### Syntax

```
default-actions {  
    interface-down;  
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

Define the action to be taken when connectivity to the remote MEP is lost.

### Default

If no action is configured, no action is taken.

### Options

**interface-down**—When a remote MEP connectivity failure is detected, bring the interface down.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

| [Configuring a CFM Action Profile to Specify CFM Actions for CFM Events](#) | 589

## delay

### Syntax

```
delay delay-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles  
  profile-name calculation-weight]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

### Description

Configure the calculation weight for delay.

### Options

**delay-value**—Calculation weight for delay.

**NOTE:** This option is applicable only for two-way delay measurement.

**Range:** 1 through 65,535

**Default:** 1

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Configuring an Iterator Profile | 782](#)

[Configuring an Iterator Profile on a Switch \(CLI Procedure\)](#)

[calculation-weight | 1008](#)

## delegate-server-processing

### Syntax

```
delegate-server-processing;
```

### Hierarchy Level

[edit protocols [oam ethernet connectivity-fault-management performance-monitoring](#)]

### Release Information

Statement introduced in Junos OS Release 11.1.

### Description

For Ethernet interfaces on MX Series routers , enable server-side processing for two-way delay measurement and loss measurement.

By default, the processing is done by the Routing Engine.

### Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

## RELATED DOCUMENTATION

| [Ethernet Frame Delay Measurements Overview](#) | 690

## delay-variation

### Syntax

```
delay-variation delay-variation-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles  
  profile-name calculation-weight]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

### Description

Configure the calculation weight for delay variation.

### Options

***delay-variation-value***—Calculation weight for delay variation.

**NOTE:** This option is applicable only for two-way delay measurement.

**Range:** 1 through 65,535

**Default:** 1

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Configuring an Iterator Profile | 782](#)

[Configuring an Iterator Profile on a Switch \(CLI Procedure\)](#)

[calculation-weight | 1008](#)

## detect-loc

### Syntax

```
detect-loc;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
maintenance-association ma-name mep mep-id remote-mep mep-id ]
```

```
[edit protocols oam ethernet link-fault-management interface interface-name ]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Specify whether Ethernet OAM continuity checks are performed for an individual remote maintenance end point (MEP).

When you configure the **detect-loc** statement at [edit protocols oam ethernet link-fault-management interface *interface-name*] hierarchy level, a loss-of-continuity (LOC) defect is raised when the peer is not found within a period that is equal to 3 times the current keepalive pdu interval. When an LOC defect is raised, a syslog error message is generated .

**NOTE:** When you configure the **detect-loc** statement at the [edit protocols oam ethernet link-fault-management interface *interface-name*] hierarchy level, any **action-profile** configured to bring down the interface is executed when an LOC defect is detected. However, the **action-profile** is not executed if you have not configured **detect-loc** statement the **detect-loc** statement at the [edit protocols oam ethernet link-fault-management interface *interface-name*] hierarchy level.

To view the current LOC status of an interface, execute the **show oam ethernet link-fault-management** command.

### Default

The MEP does not generate LOC defect messages by default.

### Required Privilege Level

interface—To view this statement in the configuration.

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

#### RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564

[remote-mep](#) | 1074

## direction

### Syntax

```
direction (up | down);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

Configure the direction of the MEP.

### Options

**up**—An UP MEP CCM is transmitted out of every logical interface which is part of the same bridging or vpls instance except for the interface configured on this MEP.

**NOTE:** The **up** direction for MEP is not supported on T Series routers.

**down**—Down MEP CCMs are transmitted only out the interface configured on this MEP.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages | 564](#)

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)



## enhanced-cfm-mode

### Syntax

```
enhanced-cfm-mode enhanced-cfm-mode;
```

### Hierarchy Level

```
[edit logical-systems name protocols oam ethernet connectivity-fault-management],  
[edit protocols oam ethernet connectivity-fault-management]
```

### Release Information

Statement introduced in Junos OS Release 17.2R1 for MX Series routers.

### Description

Enables enhanced CFM mode. When you enable enhanced CFM mode, Junos OS supports 32,000 maintenance association end points (MEPs) and maintenance intermediate points (MIPs) each per chassis for bridge, VPLS, L2VPN, and CCC domains. To support enhanced CFM mode, configure the network services mode on the router as **enhanced-ip**.

**NOTE:** After enabling CFM mode, restart CFM for the changes to take effect. If you do not restart CFM, CFM automatically restarts after 1 minute.

### Required Privilege Level

Routing

### RELATED DOCUMENTATION

[Enabling Enhanced Connectivity Fault Management Mode](#) | 596

## erroneous-ccm

### Syntax

```
erroneous-ccm;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event  
  ais-trigger-condition]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the defect condition that raises an alarm indication signal when any cross-connect continuity check messages (CCMs) with an unexpected MEP ID or an erroneous maintenance domain level are received by the MEP.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## event (CFM)

### Syntax

```
event {
  adjacency-loss;
  interface-status-tlv [lower-layer-down down];
  port-status-tlv blocked;
  rdi;
}
```

### Hierarchy Level

[edit protocols [oam ethernet connectivity-fault-management action-profile](#)]

### Release Information

Statement introduced in Junos OS Release 10.1

### Description

Configure threshold values for connectivity fault management events in an action profile.

### Options

**adjacency-loss**—Connectivity is lost.

**interface-status-tlv [ lower-layer-down down ]**—Values that need to be monitored in interface status TLV.

**port-status-tlv**—Values that need to be monitored in port status TLV.

**rdi**—RDI received from some MEP.

### Required Privilege Level

**interface**—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a CFM Action Profile to Specify CFM Actions for CFM Events | 589](#)

[interface-status-tlv | 1036](#)

[port-status-tlv | 1065](#)

## flap-trap-monitor

### Syntax

```
flap-trap-monitor seconds;
```

### Hierarchy Level

```
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-fd-twoway-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-fd-twoway-threshold]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-ifdv-twoway-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-ifdv-twoway-threshold]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-flr-forward-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-flr-forward-threshold ]
[edit logical-systems name protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-flr-backward-threshold],
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles profile-name avg-flr-backward-threshold]
```

### Release Information

Statement introduced in Junos OS Release 17.2R1 for MX Series routers.

### Description

Enables damping of jnxSoamPmThresholdCrossingAlarm traps sent to the network management system (NMS) by summarizing the flap occurrences over a period of time and sends a single jnxSoamPmThresholdFlapAlarm notification to the NMS. You can enable damping at the global level for the iterator or you can enable damping at the individual threshold type of an iterator. You can specify the duration of time for summarizing flap occurrences.

### Options

**seconds**—Duration in seconds, for summarizing flap occurrences and send out a single jnxSoamPmThresholdFlapAlarm notification to the NMS.

**Range:** 1 through 360 seconds

**Required Privilege Level**  
routing

RELATED DOCUMENTATION

|                                                                                                                                   |
|-----------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">avg-fd-twoway-threshold</a>   <a href="#">1000</a>                                                                    |
| <a href="#">avg-ifdv-twoway-threshold</a>   <a href="#">1002</a>                                                                  |
| <a href="#">avg-flr-backward-threshold</a>   <a href="#">1006</a>                                                                 |
| <a href="#">avg-flr-forward-threshold</a>   <a href="#">1004</a>                                                                  |
| <a href="#">sla-iterator-profiles</a>   <a href="#">1079</a>                                                                      |
| <a href="#">Damping CFM performance Monitoring Traps and Notifications to Prevent Congestion of The NMS</a>   <a href="#">605</a> |

## hardware-assisted-timestamping

### Syntax

```
hardware-assisted-timestamping;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
```

### Release Information

Statement introduced in Junos OS Release 9.5.

### Description

For Ethernet interfaces on Enhanced and Enhanced Queuing Dense Port Concentrators (DPCs) in MX Series routers only, enable hardware-assisted timestamping support for Ethernet frame delay measurement.

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Frame Delay Measurements Overview | 690](#)

[Guidelines for Configuring Routers to Support an ETH-DM Session | 710](#)

[Enabling the Hardware-Assisted Timestamping Option | 722](#)

## hardware-assisted-keepalives

### Syntax

```
hardware-assisted-keepalives [enable | disable];
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
```

### Release Information

Statement introduced in Junos OS Release 14.2R1.

### Description

For Ethernet interfaces on Modular Port Concentrators (MPCs) in MX Series routers only, delegate the transmission of the continuity check messages (CCMs) to the forwarding ASIC (that is, to the hardware) by enabling inline transmission of CCMs. Inline transmission of CCMs is also known as Inline-KA.

By default, CCMs are transmitted by the CPU of the MPC and not by the hardware. If the duration between transmissions of CCMs is low or if the CCMs for a specific line card scale, we recommend that you enable delegation of the transmission of CCMs to the hardware. By enabling inline transmission of CCMs, you can achieve maximum scaling of CCMs.

**NOTE:** Starting in Junos OS Release 16.1R1, the inline support for 1s interval is supported on MPC1 to MPC7 (except MPC3 and MPC4) line cards.

### Default

Inline transmission is disabled by default.

### Options

**enable**—Enable inline transmission of CCMs.

**NOTE:** Inline transmission of CCMs is not enabled when there is a CFM session already established. To enable inline transmission, you must first deactivate the CFM session using the **deactivate** command and then reactivate the CFM session using the **activate** command.

**disable**—Disable inline transmission of CCMs.

**NOTE:** After disabling inline transmission of CCMs, you must reboot the router for the changes to take effect.

**Required Privilege Level**

interface—To view this statement in the configuration.

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

**RELATED DOCUMENTATION**

[Enabling Inline Transmission of Continuity Check Messages for Maximum Scaling | 825](#)

[Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades | 599](#)



## hold-interval (OAM)

### Syntax

```
hold-interval minutes;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name continuity-check]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

### Description

The time to wait in minutes before flushing the maintenance association end point (MEP) database, if no updates occur. The configurable range is 1 minute through 30240 minutes. The default value is 10 minutes.

**NOTE:** Hold timer based flushing is applicable only for auto discovered remote MEPs and not for statically configured remote MEPs.

### Options

***minutes***—Time to wait, in minutes.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Continuity Check Protocol Parameters Overview | 562](#)

[Configuring Continuity Check Protocol Parameters for Fault Detection | 563](#)

## instance

### Syntax

```
instance vpls-instance-name;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain name]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

### Description

Specify the VPLS instance of the default maintenance domain.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring Maintenance Intermediate Points \(MIPs\) | 556](#)

[maintenance-domain | 1051](#)

## interface (IEEE 802.1ag OAM Connectivity-Fault Management)

### Syntax

```
interface (interface-name | ((ge- | xe-) (fpc/pic/port | fpc/pic/port.unit-number | fpc/pic/port.unit-number vlan vlan-id)));
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
maintenance-association ma-name mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.1ag Operation, Administration, and Management (OAM) support.

For Gigabit Ethernet interfaces and 10-Gigabit Ethernet interfaces on MX Series routers, configure IEEE 802.1ag Connectivity Fault Management (CFM) support on trunk interface ports.

Starting in Junos OS 17.4R1, you can enable support for IEEE 802.1ag CFM on pseudowire service interfaces by configuring maintenance intermediate points (MIPs) on the pseudowire service interfaces.

**NOTE:** The CFM MIP session is supported only on the pseudowire services interface and not on the pseudowire services tunnel interface.

### Options

**interface-name**—Interface to which the MEP is attached. It could be a physical Ethernet interface, logical Ethernet interface, pseudowire services interfaces, or on a specific VLAN of a trunk port interface (MX Series only).

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564

## interface (OAM Link-Fault Management)

### 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**]

### Release Information

Statement introduced in Junos OS Release 8.2.

### 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.

## RELATED DOCUMENTATION

| [Enabling IEEE 802.3ah OAM Support](#)

## interface-down

### Syntax

```
interface-down;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name default-actions]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### Description

Bring the interface down when a remote MEP connectivity failure is detected.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

| [Configuring a CFM Action Profile to Specify CFM Actions for CFM Events](#) | 589

## interface-status-tlv

### Syntax

```
interface-status-tlv [ down lower-layer-down ];
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile profile-name event]
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
 maintenance-association ma-name continuity-check]
```

### Release Information

Statement introduced in Junos OS Release 9.6.

### Description

Defines an **action-profile** consisting of various events and the action. Based on values of **interface-status-tlv** in the received CCM packets, specific action such as *interface-down* can be taken using **action-profile** options.

### Options

**down**—When the incoming CCM packet contains interface status TLV with value down, the action will be triggered for this action-profile.

**lower-layer-down**—When the incoming CCM packet contains interface status TLV with value lower-layer-down, the action will be triggered for this action-profile.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring Port Status TLV and Interface Status TLV](#) | 626

## interface-status-send-rdi

### Syntax

```
interface-status-send-rdi interface-status-send-rdi;
```

### Hierarchy Level

```
[edit logical-systems name protocols oam ethernet connectivity-fault-management maintenance-domain name
maintenance-association name continuity-check],
[edit protocols oam ethernet connectivity-fault-management maintenance-domain name maintenance-association
name continuity-check]
```

### Release Information

Statement introduced in Junos OS Release 17.3R1 for MX Series Routers.

### Description

Configure CFM to propagate the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages when the interface is down. When the status of the EVPN provider edge device is standby, the EVPN VPWS service is notified and it sets the interface status to CCC-down. When the interface status is CCC-down, it indicates that the provider edge service is down. When you enable CFM monitoring, CFM propagates the status of the provider edge device via the remote defect indication (RDI) bit in the CC messages. Thus, the customer edge device is aware that the provider edge device is down.

Usually, when the interface goes down, CFM propagates the status of the provider edge device via the interface status TLV. If the customer edge device does not support the interface status TLV, you can use the RDI bit to propagate the status of the provider edge device.

### Required Privilege Level

routing—To view this statement in the configuration.

### RELATED DOCUMENTATION

[interface-status-tlv | 1036](#)

[Configuring Port Status TLV and Interface Status TLV | 626](#)

[Understanding CFM Monitoring between CE and PE Devices | 625](#)

## interval

### Syntax

```
interval (100ms | 10m | 10ms | 10s | 1m | 1s| 3.3ms);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name
  maintenance-association ma-name continuity-check]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

Option **10ms** introduced in Junos OS Release 9.1.

Third-party interoperability during a unified in-service software upgrade (ISSU) introduced in Junos OS Release 17.1.

Option **3.3ms** introduced in Junos OS Release 17.2.

### Description

Configure the interval between successive transmissions of continuity check messages (CCMs) as part of the connectivity fault detection strategy. When the receiving maintenance association end point (MEP) does not receive a CCM at the configured interval, the **loss-threshold** statement determines how many CCMs can be lost before the sending MEP is marked as down. The **hold-interval** statement then determines the frequency at which the database of MEPs in the maintenance association (MA) is flushed in the absence of updates.

**NOTE:** When the continuity-check interval is set to 3.3ms, the actual interval is 4ms (250 pps).

During a unified in-service software upgrade (ISSU), Junos OS connectivity fault management (CFM) works when the peer device is not a Juniper Networks router. Interoperating with the router of another vendor, the Juniper Networks router retains session information and continues to transmit CCM (continuity check message) PDUs during the unified ISSU upgrade. For this feature to work, you must enable Packet Forwarding Engine keepalives with the **hardware-assisted-keepalives** statement, and configure the interval between CCMs to be 1 second with **interval** statement.



**NOTE:** For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine by default. You can disable PPM only on the Packet Forwarding Engine. To disable PPM on the Packet Forwarding Engine, use the **no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level.

**NOTE:** A continuity check interval of 10 milliseconds is not supported for CFM sessions over a label-switched interface (LSI).

### Options

**100ms**—100 milliseconds.

**10m**—10 minutes.

**10ms**—10 milliseconds.

**10s**—10 seconds.

**1m**—1 minute.

**1s**—1 second.

**3.3ms**—3.3 milliseconds.

**Default:** 1m

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Continuity Check Protocol Parameters Overview | 562](#)

[Configuring Continuity Check Protocol Parameters for Fault Detection | 563](#)

[Configuring Connectivity Fault Management for Interoperability During Unified In-Service Software Upgrades | 599](#)

## interval (CFM MEP)

### Syntax

```
interval (1m | 1s);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action  
  log-and-generate-ais]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the interval between AIS messages that are to be received by the MEP as either 1 minute or 1 second.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## iteration-count

### Syntax

```
iteration-count count-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name  
  maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile profile-name]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

### Description

Configure the number of iterations for which the connection partakes in the iterator for acquiring SLA measurements.

### Options

***count-value***—Number of iterations for which the connection should partake in the iterator for acquiring SLA measurements.

**Range:** 1 through 65,535

**Default:** 0 (or infinite iterations)

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

### RELATED DOCUMENTATION

[sla-iterator-profile](#) | **1077**

[Configuring a Remote MEP with an Iterator Profile](#) | **794**

## iteration-period

### Syntax

```
iteration-period iteration-period-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles  
  profile-name]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

### Description

Configure the iteration period, which is the maximum number of cycles per iteration (that is, the number of connections registered to an iterator cannot exceed this value).

### Options

***iteration-period-value***—Maximum number of cycles per iteration.

**Range:** 1 through 2000

**Default:** 2000

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Configuring an Iterator Profile | 782](#)

*Configuring an Iterator Profile on a Switch (CLI Procedure)*

## level

### Syntax

```
level number;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in junos os release 12.1X48 for PTX Series Packet Transport Routers.

### Description

A number used in connectivity fault management (CFM) messages to identify the maintenance association. The number is embedded in each of the CFM frames. CFM messages within a given level are processed by maintenance end points (MEPs) at the same level. For example, the operator domain can be level 0, the provider domain can be level 3, and the customer domain can be level 7.

### Options

***number***—A number used to identify the maintenance domain to which the CFM message belongs.

**Range:** 0 through 7

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

| [Creating a Maintenance Domain](#) | 555

## level (CFM MEP)

### Syntax

```
level value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action  
  log-and-generate-ais]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the server maintenance domain level for the MEP.

### Options

**value**—Maintenance domain level.

**Range:** 1 through 7

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## linktrace

### Syntax

```
linktrace {  
  age (30m | 10m | 1m | 30s | 10s);  
  path-database-size path-database-size;  
}
```

### Hierarchy Level

[edit protocols [oam ethernet connectivity-fault-management](#)]

### Release Information

Statement introduced in Junos OS Release 8.5.

### Description

Configure connectivity fault management linktrace parameters.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring Linktrace Protocol in CFM](#) | 576

## log-and-generate-ais

### Syntax

```
log-and-generate-ais {  
  interval (1m | 1s);  
  level level;  
  priority level;  
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the action be taken when an AIS alarm is detected. The action includes generating and logging the AIS statistics along with the interval between AIS messages, the server maintenance domain level, and the priority of the AIS message.

### Options

The other statements are explained separately.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)



## loss-threshold

### Syntax

```
loss-threshold number;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name continuity-check]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

### Description

Specify the number of continuity check messages lost before marking the remote MEP as down. The value can be from 3 to 256 protocol data units (PDUs). The default value is 3 PDUs.

### Options

***number***—The number of continuity check messages that can be lost before the remote MEP is considered down.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Continuity Check Protocol Parameters Overview](#) | 562

[Configuring Continuity Check Protocol Parameters for Fault Detection](#) | 563

## lowest-priority-defect

### Syntax

```
lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect | rem-err-xcon | xcon)
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 10.0.

### Description

Specify the lowest priority defect that is allowed to generate a Fault Alarm whenever CFM detects a defect. This configuration is done at the MEP level.

### Options

Specify one of the following lowest priority defect options:

**all-defects**—Allows all defects.

**err-xcon**—Allows only erroneous CCM and cross-connect CCM defects.

**mac-rem-err-xcon**—Allows only MAC, not receiving CCM, erroneous CCM, and cross-connect defects.

**no-defect**—Allows no defects.

**rem-err-xcon**—Allows only not receiving CCM, erroneous CCM, and cross-connect CCM defects.

**xcon**—Allows only cross-connect CCM defects.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564

## maintenance-association

### Syntax

```

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 {
    hold-interval minutes;
    interval (10m | 10s | 1m | 1s | 100ms);
    loss-threshold number;
  }
  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;
      }
    }
  }
}

```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

### Description

Configure the name of the maintenance association in IEEE-compliant format.

### Options

**ma-name**—The name of the maintenance association within the maintenance domain.

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

**Required Privilege Level**

interface—To view this statement in the configuration.

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

**RELATED DOCUMENTATION**

---

[Creating a Maintenance Association | 561](#)

[Configuring a MEP to Generate and Respond to CFM Protocol Messages | 564](#)

## maintenance-domain

### Syntax

```

maintenance-domain domain-name {
  bridge-domain name <vlan-id [ vlan-ids ]>;
  instance vpls-instance-name;
  level number;
  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 {
      hold-interval minutes;
      interval (10m | 10s | 1m | 1s | 100ms);
      loss-threshold number
    }
  }
  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;
      }
    }
  }
  mip-half-function(none | default | explicit);
  name-format (character-string | none | dns | mac+2oct);
}
virtual-switch name {
  bridge-domain name <vlan-id [ vlan-ids ]>;
}
}

```

### Hierarchy Level

[edit protocols [oam ethernet connectivity-fault-management](#)]

### Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Support for multiple down MEP introduced in Junos OS Release 15.1R1 for MX Series Routers.

### Description

Configure the name of the maintenance domain in IEEE-compliant format.

**NOTE:** For MX Series Routers, you can configure multiple down MEPs for a single instance of maintenance domain identifier and maintenance association name to monitor services provided on Virtual Private LAN Service (VPLS), bridge, circuit cross-connect (CCC), and IPv4 domains.

### Options

***domain-name***—Name of the maintenance domain.

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Creating a Maintenance Domain | 555](#)

[Configuring a MEP to Generate and Respond to CFM Protocol Messages | 564](#)

## measurement-interval

### Syntax

measurement-interval (5|15|30|60)

### Hierarchy Level

[edit protocols oam ethernet cfm performance-monitoring]

[edit protocols oam ethernet cfm performance-monitoring sla-iterator-profiles *profile-name*]

### Release Information

### Description

Configure measurement interval to be used for a performance monitoring session. You must configure the measurement-interval at the [edit protocols oam ethernet cfm performance-monitoring] hierarchy level, which is a global level parameter. You can override the configured value by specifying a measurement-interval for the iterator profile at the [edit protocols oam ethernet cfm performance-monitoring sla-iterator-profiles *profile-name*] hierarchy level.

**NOTE:** When you configure when MEF-36-compliant performance monitoring, you must also configure an **enhanced-sla-iterator** at the [edit protocols oam ethernet cfm performance-monitoring] hierarchy level.

### Default

15 minutes

### Required Privilege Level

configure—To enter configuration mode.

control—To modify any configuration.

## RELATED DOCUMENTATION

[IEEE 802.1ag OAM Connectivity Fault Management Overview | 548](#)

[Junos OS Support for Performance Monitoring Compliant with Technical Specification MEF 36 | 604](#)

## measurement-type

### Syntax

```
measurement-type (loss | statistical-loss-measurement | two-way-delay);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles  
  profile-name]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

The **statistical-loss-measurement** option introduced in Junos OS Release 11.2.

### Description

Configure the measurement type for the service level agreement (SLA) frames. An SLA frame is a type of packet used to measure frame loss in Ethernet connections.

### Options

**loss**—Use Y.1731-compliant line module (LM) frames to measure frame loss.

**statistical-loss-measurement**— Use Y.1731-compliant two-way data module (DM) frames to statistically measure frame loss.

**two-way-delay**—Use Y.1731-compliant two-way DM frames to measure frame loss.

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

### RELATED DOCUMENTATION

| [Configuring an Iterator Profile](#) | 782



## mep

### Syntax

```
mep mep-id {
  action-profile action-profile-name
  auto-discovery;
  direction (up | down);
  interface interface-name (protect | working);
  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;
    }
  }
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

The numeric identifier of the maintenance association end point (MEP) within the maintenance association.

### Options

**mep mep-id**—Specify the numeric identifier of the MEP.

**Range:** 1 through 8191

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

Ethernet Alarm Indication Signal (ETH-AIS) Function Overview | **810**

---

Configuring ETH-AIS on a CFM MEP | **816**

---

Configuring a MEP to Generate and Respond to CFM Protocol Messages | **564**

## mip-half-function

### Syntax

```
mip-half-function (none | default | explicit);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-managementmaintenance-domain md-name],  
[edit protocols oam ethernet connectivity-fault-managementmaintenance-association ma-name]
```

### Release Information

Statement introduced in Junos OS Release 9.6.

### Description

Specify the OAM Ethernet CFM maintenance domain MIP half functions.

**NOTE:** Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the **mip-half-function** value for all maintenance domains and maintenance associations are the same.

### Options

**none**—Specify to not use the mip-half-function.

**default**—Specify to use the default mip-half-function.

**explicit**—Specify an explicit mip-half-function.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Creating a Maintenance Domain | 555](#)

[maintenance-domain | 1051](#)

## name-format

### Syntax

```
name-format (character-string | none | dns | mac+2oct);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

### Description

Specify the format of the maintenance domain name.

### Options

**character-string**—The name is an ASCII character string.

**none**—The maintenance domain name is not used.

**dns**—The name is in domain name service (DNS) format. For example: www.juniper.net.

**mac+2oct**—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100.

**Default:** character-string

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Creating a Maintenance Association | 561](#)

[Creating a Maintenance Domain | 555](#)

## no-aggregate-delegate-processing

### Syntax

```
no-aggregate-delegate-processing;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

### Release Information

Statement introduced in Junos OS Release 14.1.

### Description

Disable distribution of connectivity fault management (CFM) sessions on aggregated Ethernet interfaces.

### Default

CFM sessions on aggregated Ethernet interfaces are distributed by default.

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

| [IEEE 802.1ag OAM Connectivity Fault Management Overview](#) | 548

## path-database-size

### Syntax

```
path-database-size path-database-size;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management linktrace]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### Description

Number of linktrace reply entries to be stored per linktrace request.

### Options

**path-database-size**—Database size.

**Range:** 1 through 255

**Default:** 64

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring Linktrace Protocol in CFM](#) | 576

## performance-monitoring

### Syntax

```
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 connectivity-fault-management](#)]

### Release Information

Statement introduced in Junos OS Release 9.5.

### Description

Specify performance monitoring support for Ethernet frame delay measurement.

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

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Ethernet Frame Delay Measurements Overview | 690](#)

[Guidelines for Configuring Routers to Support an ETH-DM Session | 710](#)

[Enabling the Hardware-Assisted Timestamping Option | 722](#)



## policer (CFM Global)

### Syntax

```
policer {  
    all cfm-policer-name;  
    continuity-check cfm-policer-name;  
    other cfm-policer-name;  
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

### Release Information

Statement introduced in Junos OS Release 10.0.

### Description

Specify a policer at the global level to police the CFM traffic belonging to all sessions.

### Options

**continuity-check *cfm-policer-name***—Police all continuity check packets with the policer specified.

**other *cfm-policer-name***—Police all non-continuity check packets with the policer specified.

**all *cfm-policer-name***—Police all CFM packets with policer specified. If the **all** option is used, then you cannot specify above two options.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Configuring Rate Limiting of Ethernet OAM Messages | 576](#)

[policer \(CFM Session\) | 1064](#)

## policer (CFM Session)

### Syntax

```
policer {
  all cfm-policer-name;
  continuity-check cfm-policer-name;
  other cfm-policer-name;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain name level number
  maintenance-association name]
```

### Release Information

Statement introduced in Junos OS Release 10.0.

### Description

Specify a separate policer to rate-limit packets specific to that session.

### Options

- **continuity-check *cfm-policer-name***—Police continuity check packets belonging to this session.
- other *cfm-policer-name***—Police all non-continuity check packets belonging to this session.
- all *cfm-policer-name***—Police all CFM packets belonging to this session. If the **all** option is used, then you cannot specify the above two options.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Configuring Rate Limiting of Ethernet OAM Messages | 576](#)

[policer \(CFM Global\) | 1063](#)

## port-status-tlv

### Syntax

```
port-status-tlv blocked;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile tlv-action event]  
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
maintenance-association ma-name continuity-check]
```

### Release Information

Statement introduced in Junos OS Release 9.6.

### Description

Define an **action-profile** consisting of various events and the action. Based on values of **port-status-tlv** in the received CCM packets, specific action such as *interface-down* can be taken using **action-profile** options.

### Options

**blocked**—When the incoming CCM packet contains port status TLV with value blocked, the action will be triggered for this action-profile.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a CFM Action Profile to Specify CFM Actions for CFM Events | 589](#)

[Configuring Port Status TLV and Interface Status TLV | 626](#)

## priority (Protocols OAM)

### Syntax

```
priority priority-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id remote-mep remote-mep-id sla-iterator-profile
  profile-name sla-iterator-profile]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

### Description

Configure the priority of the iterator profile, which is the **vlan-pcp** value that is sent in the Y.1731 data frames.

### Options

**priority-value**—Priority value, which is the **vlan-pcp** value that is sent in the Y.1731 data frames.

**Range:** 0 through 7

**Default:** 0

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[sla-iterator-profile](#) | 1077

[Configuring a Remote MEP with an Iterator Profile](#) | 794

## priority (CFM MEP)

### Syntax

```
priority value;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action log-generate-ais]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the 802.1p priority of the AIS packet.

### Options

***value***—Priority level.

**Range:** 0 through 7

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## priority (OAM Connectivity-Fault Management)

### Syntax

```
priority number;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name maintenance-association
ma-name mep mep-id]
```

For EX Series Switches:

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name maintenance-association ma-name mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

IEEE 802.1p priority bits used by the continuity check messages.

### Options

***number***—Configure the IEEE 802.1p priority bits to be used in the VLAN header of the CFM packets.

**Range:** 0 through 7

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564

## protocol (Server MEP)

### Syntax

```
protocol (l2circuit | l2vpn | ethernet) {
    interface interface-name;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the protocol as Layer 2 circuit, Layer 2 VPN, or Ethernet and associate the interface to the protocol that needs to be monitored for ETH-AIS.

### Options

**l2circuit**—Configure the protocol for the server MEP as Layer 2 circuit.

**l2vpn**—Configure the protocol for the server MEP as Layer 2 VPN.

**ethernet**—Configure the protocol for the server MEP as Ethernet.

**interface *interface-name***—The interface that is to be associated with the protocol that needs monitoring for ETH-AIS.

### Required Privilege Level

**interface**—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## protect-maintenance-association (OAM)

### Syntax

```
protect-maintenance-association protect-ma-name;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name ]
```

### Release Information

Statement introduced in Junos OS Release 11.4

### Description

Configure the name of the protect transport path for the maintenance-association.

### Options

***protect-ma-name***—The name of the protect transport path.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564



## rdi

### Syntax

```
rdi;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile tlv-action event]
```

### Release Information

Statement introduced in Junos OS Release 10.1.

### Description

Define a new event **rdi**. The remote defect indication (rdi) event is triggered whenever CCM packets are received from a remote location with the rdi bit set.

This event is cleared and action is reverted when none of the remote MEPs send the CCM packets with the RDI bit.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring a CFM Action Profile to Specify CFM Actions for CFM Events](#) | 589

## receive-ais

### Syntax

```
receive-ais;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event  
  ais-trigger-condition]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Send a message to the peer MEPs when an AIS message is received by a peer MEP at its own maintenance level.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## remote-maintenance-association (OAM)

### Syntax

```
remote-maintenance-association remote-ma-name;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name]
```

### Release Information

Statement introduced in Junos OS Release 11.4.

### Description

Configure the name of the remote maintenance association.

### Options

***remote-ma-name***—Name of the remote maintenance association.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564

## remote-mep

### Syntax

```
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;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

Configure the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.

### Options

**mep-id**—Numeric identifier of the MEP.

**Range:** 1 through 8191

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

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Configuring a MEP to Generate and Respond to CFM Protocol Messages](#) | 564

[detect-loc](#) | 1020

## sendid-tlv

### Syntax

```
sendid-tlv {  
    send-chassis-tlv;  
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management]  
[edit protocols oam ethernet connectivity-fault-management maintenance-domain  
maintenance-domain-name maintenance-association maintenance-association-name  
continuity-check]
```

### Release Information

Statement introduced in Junos OS Release 16.1R2.

### Description

Configures Junos OS to send the sender ID TLV along with the packets. The sender ID TLV is an optional TLV that is sent in continuity check messages (CCMs), loopback messages, and Link Trace Messages (LTMs), as specified in the IEEE 802.1ag standard. TLVs (type, length, and value) are described in the IEEE 802.1ag standard for Connectivity Fault Management (CFM) as a method of encoding variable-length and optional information in a protocol data unit (PDU).

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring Chassis ID TLV](#) | 640

## short-name-format

### Syntax

```
short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name  
  maintenance-association ma-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

### Description

Specify the name format of the maintenance association name.

### Options

**character-string**—The name is an ASCII character string.

**vlan**—The primary VLAN identifier.

**2octet**—A number in the range 0 through 65,535.

**rfc-2685-vpn-id**—A VPN identifier that complies with RFC 2685.

**Default:** character-string

**NOTE:** The PTX Series Packet Transport Routers support the **vlan** and **2octet** options only.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Creating a Maintenance Association](#) | 561

## sla-iterator-profile

### Syntax

```
sla-iterator-profile profile-name {
  data-tlv-size size;
  iteration-count count-value;
  priority priority-value;
}
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain md-name
  maintenance-association ma-name mep mep-id remote-mep remote-mep-id]
```

### Release Information

Statement introduced in Junos OS Release 11.1.

### Description

Configure a remote MEP with an iterator profile and specify the options.

Y.1731 performance monitoring (PM) over Aggregated Ethernet Interfaces is not supported on EX4300 switches.

### Options

**profile-name**—Name of the iterator profile configured for a remote MEP. For more information about configuring a remote MEP with an iterator profile, see [“Configuring a Remote MEP with an Iterator Profile” on page 794](#).

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

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

## RELATED DOCUMENTATION

[Configuring an Iterator Profile | 782](#)

[Configuring a Remote MEP with an Iterator Profile | 794](#)

[Verifying the Configuration of an Iterator Profile | 785](#)

[Managing Iterator Statistics | 788](#)





## sla-iterator-profiles

### Syntax

```
sla-iterator-profiles {
  profile-name {
    avg-fd-twoway-threshold;
    avg-ifdv-twoway-threshold;
    avg-flr-forward-threshold;
    avg-flr-backward-threshold;
    calculation-weight {
      delay delay-weight;
      delay-variation delay-variation-weight;
    }
    cycle-time milliseconds;
    flap-trap-monitor seconds
    iteration-period iteration-period-value;
    measurement-type (loss | statistical-frame-loss | two-way-delay);
  }
}
```

### Hierarchy Level

[edit protocols [oam ethernet connectivity-fault-management performance-monitoring](#)]

### Release Information

Statement introduced in Junos OS Release 11.1.

### Description

Configure an iterator application and specify the iterator profile options.

### Options

**profile-name**—Name of the iterator profile. For more information about configuring the iterator profile, see [“Configuring an Iterator Profile” on page 782](#).

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

### Required Privilege Level

Configure—To enter configuration mode.

Control—To modify any configuration.

### RELATED DOCUMENTATION

[Configuring an Iterator Profile | 782](#)

---

[Configuring a Remote MEP with an Iterator Profile | 794](#)

---

[Verifying the Configuration of an Iterator Profile | 785](#)

---

[Managing Iterator Statistics | 788](#)

# Configuration Statements

## IN THIS CHAPTER

- `accept-source-mac` | 1088
- `access-concentrator` | 1090
- `account-layer2-overhead` (PIC Level) | 1091
- `action` (OAM) | 1092
- `action-profile` | 1093
- `adaptive` | 1095
- `address` | 1097
- `adjacency-loss` | 1101
- `age` | 1102
- `agent-specifier` | 1103
- `aggregate` (Gigabit Ethernet CoS Policer) | 1105
- `aggregated-devices` | 1106
- `aggregated-ether-options` | 1108
- `alarms` | 1110
- `allow-remote-loopback` | 1111
- `apply-action-profile` | 1112
- `asynchronous-notification` | 1113
- `auto-negotiation` | 1114
- `auto-reconnect` | 1116
- `bandwidth-limit` (Policer for Gigabit Ethernet Interfaces) | 1117
- `bert-algorithm` | 1118
- `bert-error-rate` | 1121
- `bert-period` | 1123
- `bridge-domain` | 1125
- `burst-size-limit` (Policer for Gigabit Ethernet Interfaces) | 1126
- `centralized` | 1127
- `cfp-to-et` | 1128
- `classifier` | 1129

- client | **1130**
- community-vlans (MX Series) | **1131**
- connectivity-fault-management | **1132**
- delay (PPPoE Service Name Tables) | **1136**
- destination (IPCP) | **1137**
- device-count | **1138**
- disable (Link Protection) | **1139**
- distribution-list | **1140**
- drop (PPPoE Service Name Tables) | **1141**
- dynamic-profile (PPPoE Service Name Tables) | **1142**
- egress-policer-overhead | **1143**
- encapsulation (Logical Interface) | **1145**
- encapsulation | **1150**
- ether-options | **1157**
- ethernet (Chassis) | **1159**
- ethernet (Protocols OAM) | **1160**
- ethernet-policer-profile | **1168**
- evcs | **1170**
- evc-protocol cfm | **1171**
- event (LFM) | **1173**
- event-thresholds | **1175**
- family | **1176**
- fastether-options | **1182**
- flow-control | **1184**
- fnp | **1186**
- force-up | **1187**
- forwarding-class (Gigabit Ethernet IQ Classifier) | **1188**
- forwarding-mode (100-Gigabit Ethernet) | **1189**
- forwarding-mode (PTX Series Packet Transport Routers) | **1190**
- frame-error | **1191**
- frame-period | **1192**
- frame-period-summary | **1193**
- framing (10-Gigabit Ethernet Interfaces) | **1194**
- ggether-options | **1196**

- [hold-time up](#) | 1198
- [ieee802.1p](#) | 1199
- [ignore-l3-incompletes](#) | 1200
- [ingress-policer-overhead](#) | 1201
- [ingress-rate-limit](#) | 1203
- [inline](#) | 1204
- [input-policer](#) | 1206
- [input-priority-map](#) | 1207
- [input-three-color](#) | 1208
- [input-vlan-map \(Aggregated Ethernet\)](#) | 1209
- [interface](#) | 1210
- [interface \(OAM Link-Fault Management\)](#) | 1212
- [interface-group](#) | 1214
- [interface-group-down](#) | 1215
- [interface-none](#) | 1216
- [isolated-vlan \(MX Series\)](#) | 1217
- [lACP \(802.3ad\)](#) | 1218
- [lACP \(Aggregated Ethernet\)](#) | 1219
- [lACP \(Protocols\)](#) | 1222
- [layer2-policer](#) | 1224
- [link-adjacency-loss](#) | 1225
- [link-discovery](#) | 1226
- [link-degrade-monitor](#) | 1227
- [link-down](#) | 1228
- [link-event-rate](#) | 1229
- [link-fault-management](#) | 1230
- [link-mode](#) | 1232
- [link-protection](#) | 1234
- [link-protection \(non-LACP\)](#) | 1236
- [link-speed \(Aggregated Ethernet\)](#) | 1238
- [link-speed \(Aggregated SONET/SDH\)](#) | 1240
- [lmi \(Ethernet OAM\)](#) | 1241
- [load-balance](#) | 1243
- [load-balance-stateful \(Aggregated Ethernet Interfaces\)](#) | 1245

- load-type (Aggregated Ethernet Interfaces) | 1246
- local-bias (ae load-balance) | 1247
- logical-tunnel-options | 1249
- loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | 1251
- loopback (Local and Remote) | 1253
- loopback-tracking | 1254
- loss-priority | 1255
- mac | 1256
- mac-address (Accept Source Mac) | 1257
- mac-learn-enable | 1258
- mac-validate | 1259
- master-only | 1260
- max-sessions (PPPoE Service Name Tables) | 1261
- max-sessions-vsa-ignore (Static and Dynamic Subscribers) | 1262
- maximum-links | 1264
- mc-ae | 1266
- minimum-bandwidth (aggregated Ethernet) | 1272
- minimum-links | 1273
- mixed-rate-mode | 1274
- mtu | 1275
- mru | 1279
- multicast-statistics | 1280
- negotiate-address | 1281
- negotiation-options | 1282
- no-adaptive | 1283
- no-allow-link-events | 1284
- no-auto-mdix | 1285
- no-keepalives | 1286
- no-pre-classifier | 1287
- no-send-pads-ac-info | 1288
- no-send-pads-error | 1289
- non-revertive (Interfaces) | 1290
- number-of-sub-ports | 1291
- oam | 1293

- optics-options | 1297
- otn-options | 1299
- output-policer | 1301
- output-priority-map | 1302
- output-three-color | 1303
- output-vlan-map (Aggregated Ethernet) | 1304
- pado-advertise | 1305
- passive-monitor-mode | 1306
- pdu-interval | 1307
- pdu-threshold | 1308
- per-flow (Aggregated Ethernet Interfaces) | 1309
- periodic | 1310
- policer (CFM Firewall) | 1312
- policer (CoS) | 1313
- policer (MAC) | 1315
- port-priority | 1317
- pp0 (Dynamic PPPoE) | 1318
- ppm (Ethernet Switching) | 1320
- pppoe-options | 1321
- pppoe-underlying-options (Static and Dynamic Subscribers) | 1323
- preferred-source-address | 1324
- premium (Output Priority Map) | 1326
- premium (Policer) | 1327
- protocol-down | 1328
- premium (Output Priority Map) | 1329
- premium (Policer) | 1330
- proxy | 1331
- rebalance (Aggregated Ethernet Interfaces) | 1332
- receive-options-packets | 1333
- receive-ttl-exceeded | 1334
- recovery | 1335
- remote-loopback | 1337
- restore-interval | 1338
- revertive | 1339

- [routing-instance](#) | **1340**
- [routing-instance \(PPPoE Service Name Tables\)](#) | **1341**
- [rx-enable](#) | **1342**
- [rx-max-duration](#) | **1343**
- [sa-multicast \(100-Gigabit Ethernet\)](#) | **1344**
- [sa-multicast \(PTX Series Packet Transport Routers\)](#) | **1346**
- [send-critical-event](#) | **1347**
- [server](#) | **1348**
- [service \(PPPoE\)](#) | **1349**
- [service-name](#) | **1351**
- [service-name-table](#) | **1352**
- [service-name-tables](#) | **1354**
- [session-expiry \(MX Series in Enhanced LAN Mode\)](#) | **1355**
- [source-address-filter](#) | **1356**
- [source-filtering](#) | **1358**
- [speed \(Ethernet\)](#) | **1359**
- [speed \(MX Series DPC\)](#) | **1368**
- [static-interface](#) | **1369**
- [switch-options](#) | **1370**
- [switch-port](#) | **1371**
- [symbol-period](#) | **1372**
- [sync-reset](#) | **1373**
- [syslog \(OAM Action\)](#) | **1374**
- [system-id](#) | **1375**
- [system-priority](#) | **1376**
- [targeted-options \(Grouping Subscribers by Bandwidth Usage\)](#) | **1377**
- [targeted-options \(Manual Targeting\)](#) | **1379**
- [targeted-distribution](#) | **1381**
- [targeted-options](#) | **1382**
- [terminate \(PPPoE Service Name Tables\)](#) | **1383**
- [thresholds](#) | **1384**
- [traceoptions \(LLDP\)](#) | **1386**
- [traceoptions \(Individual Interfaces\)](#) | **1389**
- [traceoptions \(LACP\)](#) | **1398**



- [traceoptions \(PPPoE\) | 1400](#)
- [tx-duration | 1403](#)
- [tx-enable | 1404](#)
- [underlying-interface | 1405](#)
- [unit | 1407](#)
- [unnumbered-address \(Dynamic Profiles\) | 1418](#)
- [unnumbered-address \(PPP\) | 1420](#)
- [Unused | 1421](#)
- [virtual-control-channel | 1422](#)
- [virtual-switch | 1423](#)
- [vlan-rule \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1424](#)
- [vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1425](#)

## accept-source-mac

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

### Description

For Gigabit Ethernet intelligent queuing (IQ) interfaces only, 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:** 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.

#### RELATED DOCUMENTATION

---

[Configuring Gigabit Ethernet Policers | 293](#)

---

[Configuring MAC Address Filtering on PTX Series Packet Transport Routers | 23](#)

---

[source-filtering | 1358](#)

## access-concentrator

### 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]
```

### 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.

### 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.

## RELATED DOCUMENTATION

[Identifying the Access Concentrator | 49](#)

[Configuring the PPPoE Family for an Underlying Interface](#)

[Configuring Dynamic PPPoE Subscriber Interfaces](#)

[PPPoE Overview | 41](#)

## account-layer2-overhead (PIC Level)

### Syntax

```
account-layer2-overhead;
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 13.2.

### 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.

## RELATED DOCUMENTATION

[Accounting of the Layer 2 Overhead Attribute in Interface Statistics | 283](#)

[Configuring Layer 2 Overhead Accounting in Interface Statistics | 286](#)

[Verifying the Accounting of Layer 2 Overhead in Interface Statistics | 287](#)

## action (OAM)

### Syntax

```
action {  
    link-down;  
    send-critical-event;  
    syslog;  
}
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 8.5.

### Description

Define the action or actions to be taken when the OAM fault event occurs.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Specifying the Actions to Be Taken for Link-Fault Management Events](#) | 674

## action-profile

### List of Syntax

Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches on page 1093

Syntax: EX Series Switches and NFX Series Devices on page 1093

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

```
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;
  }
}
```

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

```
action-profile profile-name;
action {
  syslog;
  link-down;
}
event {
  link-adjacency-loss;
  link-event-rate {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
}
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

### Release Information

Statement introduced in Junos OS Release 8.5 for T, M, MX and ACX Series Routers, SRX Series Firewalls, and EX Series Switches, .

Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

### Description

Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.

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

### Options

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

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring an OAM Action Profile | 673](#)

*Configuring Ethernet OAM Link Fault Management*



## adaptive

### 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]
```

### Release Information

Statement introduced in Junos OS Release 13.2R3 for MX Series Routers.

Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.

Statement introduced in Junos OS Release 15.1X53-D10 for the QFX Series.

### Description

Correct a genuine traffic imbalance by using a feedback mechanism to distribute the traffic across the links of an aggregated Ethernet bundle.

### Options

**pps**—(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.

**scan-interval *multiple***—(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.

**RELATED DOCUMENTATION**

[Understanding Aggregated Ethernet Load Balancing | 141](#)

[Example: Configuring Aggregated Ethernet Load Balancing | 161](#)

## address

### Syntax

```

address address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  destination address;
  destination-profile name;
  eui-64;
  master-only;
  multipoint-destination address dlcidlcid-identifier;
  multipoint-destination address {
    epd-threshold cells;
    inverse-arp;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
      queue-length number;
    }
    vci vpi-identifier.vci-identifier;
  }
  primary;
  preferred;
  virtual-gateway-address
  (vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
  }
  priority-number number;
  track {
    priority-cost seconds;
    priority-hold-time interface-name {
      interface priority;
      bandwidth-threshold bits-per-second {
        priority;

```

```

    }
  }
  route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-address [ addresses ];
}
}

```

### Hierarchy Level

```

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

```

### Release Information

Statement introduced before Junos OS Release 7.4.

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

## Description

Configure the interface address.

**NOTE:** If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, and the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/8;
      }
    }
  }
}
```

For more information on configuring the same address on multiple interfaces, see *Configuring the Interface Address*.

- 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.

RELATED DOCUMENTATION

|                                                 |
|-------------------------------------------------|
| <i>Configuring the Protocol Family</i>          |
| <i>family</i>                                   |
| <a href="#">negotiate-address</a>   <b>1281</b> |
| <i>unnumbered-address (Ethernet)</i>            |

## adjacency-loss

### Syntax

```
adjacency-loss;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name event  
  ais-trigger-condition]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### Description

Configure the defect condition that raises an alarm indication signal when physical connectivity is lost between peer MEPs.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview | 810](#)

[Configuring ETH-AIS on a CFM MEP | 816](#)

## age

### Syntax

```
age (30m | 10m | 1m | 30s | 10s);
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management linktrace]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### 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.

## RELATED DOCUMENTATION

[Configuring Linktrace Protocol in CFM](#) | 576



## agent-specifier

### 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]
```

### 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.

### 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.

## 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)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

*premium (Hierarchical Policer)*

[ieee802.1p](#) | 1199

## aggregated-devices

### Syntax

```
aggregated-devices {
  ethernet {
    device-count number;
    lacp {
      link-protection {
        non-revertive;
      }
      system-priority;
    }
  }
  sonet {
    device-count number;
  }
  maximum-links maximum-links-limit;
}
```

### Hierarchy Level

[edit chassis]

### Release Information

Statement introduced before Junos OS Release 7.4.

Support for LACP link protection and system priority introduced in Junos OS Release 9.3.

### Description

Configure properties for aggregated devices on the router. Aggregate Ethernet links are logical interfaces defined on the device that bundle together multiple physical interfaces into a single interface for the use of redundancy and bandwidth aggregation. When interconnecting devices you can create aggregate ethernet interfaces to bundle together multiple physical ethernet links to increase bandwidth and redundancy between devices.

Link aggregation enables you to group Ethernet interfaces to form a single link layer interface. Link Aggregation Control Protocol (LACP) is supported in chassis cluster deployments, where aggregated Ethernet interfaces and redundant Ethernet interfaces are supported simultaneously.

You must first configure the system to enable configuring the Aggregated Ethernet (ae) Interfaces. By default, Juniper devices do not have any aggregated ethernet interfaces created. To configure the device to support a given number of ae interfaces, you must define it on a per chassis basis using the **set chassis aggregated-devices devices {1-32}** in configuration mode. The number of devices you define will be the number of aggregated ethernet interfaces that the system will create which can be configured just like

any other ethernet interface. Also you can view the interfaces created by using the **show interface terse** command. Once you have defined the number of aggregated ethernet devices on the chassis you can then continue to configure the LAG members on a per ethernet interface basis.

### Options

The remaining statements are explained separately.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring Junos OS for Supporting Aggregated Devices](#) | 77

## aggregated-ether-options

### 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]

## Release Information

Statement introduced before Junos OS Release 7.4.

## 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.

## RELATED DOCUMENTATION

| [Ethernet Interfaces Overview](#) | 3

## alarms

### Syntax

```
alarms;
```

### Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

### Release Information

Statement introduced in JUNOS Release 10.1.

### 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.

### RELATED DOCUMENTATION

[Ethernet DWDM Interface Wavelength Overview](#) | 505



## allow-remote-loopback

### Syntax

```
allow-remote-loopback;
```

### Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name negotiation-options]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

### RELATED DOCUMENTATION

[Enabling Remote Loopback Support on the Local Interface](#) | 681

## apply-action-profile

### Syntax

```
apply-action-profile profile-name;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### Description

Apply the specified action profile to the interface for link-fault management.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Applying an Action Profile](#) | 677

## asynchronous-notification

### Syntax

```
(asynchronous-notification | no-asynchronous-notification);
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

### 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.

## RELATED DOCUMENTATION

[Gigabit Ethernet Notification of Link Down Alarm Overview | 236](#)

[Configuring Gigabit Ethernet Notification of Link Down Alarm | 236](#)

## auto-negotiation

### 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]
```

### Release Information

Statement introduced in Junos OS Release 7.6.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

### 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. For EX Series switches, explicitly enable autonegotiation only.

- **auto-negotiation**—Enables autonegotiation. This is the default.
- **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, 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.

## 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.

## RELATED DOCUMENTATION

[Gigabit Ethernet Autonegotiation Overview | 305](#)

[Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

[Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support](#)

## auto-reconnect

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring the PPPoE Automatic Reconnect Wait Timer | 49](#)

*Junos OS Interfaces and Routing Configuration Guide*

## bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

[burst-size-limit \(Policer for Gigabit Ethernet Interfaces\)](#) | 1126

## bert-algorithm

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

### 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.

RELATED DOCUMENTATION

|                                               |
|-----------------------------------------------|
| <i>Interface Diagnostics</i>                  |
| <i>Configuring E1 BERT Properties</i>         |
| <i>Configuring E3 BERT Properties</i>         |
| <i>Configuring T1 BERT Properties</i>         |
| <i>Configuring T3 BERT Properties</i>         |
| <i>Examples: Configuring T3 Interfaces</i>    |
| <a href="#">bert-error-rate</a>   <b>1121</b> |
| <a href="#">bert-period</a>   <b>1123</b>     |

## bert-error-rate

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

### 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.

### RELATED DOCUMENTATION

[bert-algorithm](#) | 1118

[bert-period](#) | 1123

*ds0-options*

*e1-options*

*e3-options*

*t1-options*

*t3-options*

*Interface Diagnostics*

*Configuring E1 BERT Properties*

*Configuring E3 BERT Properties*

*Configuring T1 BERT Properties*

*Configuring T3 BERT Properties*

*Examples: Configuring T3 Interfaces*

## bert-period

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers.

### 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.

**RELATED DOCUMENTATION**

|                                               |
|-----------------------------------------------|
| <i>Interface Diagnostics</i>                  |
| <i>Configuring E1 BERT Properties</i>         |
| <i>Configuring E3 BERT Properties</i>         |
| <i>Configuring T1 BERT Properties</i>         |
| <i>Configuring T3 BERT Properties</i>         |
| <a href="#">bert-algorithm</a>   <b>1118</b>  |
| <a href="#">bert-error-rate</a>   <b>1121</b> |

## bridge-domain

### 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]
```

### Release Information

Statement introduced in Junos OS Release 9.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Maintenance Intermediate Points \(MIPs\) | 556](#)

[maintenance-domain | 1051](#)

## burst-size-limit (Policer for Gigabit Ethernet Interfaces)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

### RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 293](#)

[bandwidth-limit \(Policer for Gigabit Ethernet Interfaces\) | 1117](#)



## centralized

### Syntax

```
centralized;
```

### Hierarchy Level

```
[edit protocols lacp ppm]
```

### Release Information

Statement introduced in Junos OS Release 9.4 for MX Series routers.

Statement introduced in Junos OS Release 10.2 for EX Series switches.

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

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

### Description

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.

## RELATED DOCUMENTATION

[inline](#) | [1204](#)

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

---

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

---

*Configuring Distributed Periodic Packet Management*

---

*Configuring Link Aggregation*

## cfp-to-et

### Syntax

```
cfp-to-et;
```

### Hierarchy Level

```
[edit chassis fpc slot]
```

### Release Information

Statement introduced in Junos OS Release 19.2R1-S1.

### 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.

## RELATED DOCUMENTATION

---

*Interface Mapping and Modulation Format for ACX5448-D*

---

## classifier

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

### RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

[input-priority-map](#) | 1207

## client

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### 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.

### RELATED DOCUMENTATION

[Configuring the PPPoE Client Mode](#) | 50

## community-vlans (MX Series)

### 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 ],
```

### Release Information

Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers.

### 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.

### RELATED DOCUMENTATION

## connectivity-fault-management

### Syntax

```
connectivity-fault-management {
  action-profile profile-name {
    action {
      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](#)]

## Release Information

Statement introduced in Junos OS Release 8.4.

## 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](#).

## Required Privilege Level

interface—To view this statement in the configuration.

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



## RELATED DOCUMENTATION

| [IEEE 802.1ag OAM Connectivity Fault Management Overview](#) | 548

## delay (PPPoE Service Name Tables)

### 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]
```

### 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.

### 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.

### RELATED DOCUMENTATION

| [Configuring PPPoE Service Name Tables](#)

## destination (IPCP)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

*Configuring IPCP Options for Interfaces with PPP Encapsulation*

[address](#) | **1097**

[negotiate-address](#) | **1281**

*Junos OS Administration Library*

## device-count

### Syntax

```
device-count number;
```

### Hierarchy Level

```
[edit chassis aggregated-devices ethernet]
```

```
[edit chassis aggregated-devices sonet]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement functionality updated in Junos OS Release 14.2, as described below.

### 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.

## RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices | 77](#)

*Configuring Aggregated SONET/SDH Interfaces*

## disable (Link Protection)

### Syntax

```
disable;
```

### Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

### Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### Description

Disable LACP link protection on the interface.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces | 67](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## distribution-list

### Syntax

```
distribution-list distribution-list-number;
```

### Hierarchy Level

```
[edit interfaces interface name gigether-options 802.3ad ]  
[edit dynamic-profiles name interfaces name gigether-options 802.3ad]  
[edit dynamic-profiles name logical-systems name interfaces name gigether-options 802.3ad]
```

### Release Information

Statement introduced in Junos OS Release 16.1R1.

### Description

Specify a distribution list to a Gigabit Ethernet interface to carry traffic. You can then configure the distribution list as a primary list or a backup list for the members of an aggregated Ethernet bundle.

Example:

```
[edit]  
user@router# set interfaces ge-0/0/3 gigether-options 802.3ad distribution-list dl1
```

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 100](#)

[targeted-options | 1382](#)

[targeted-distribution | 1381](#)

## drop (PPPoE Service Name Tables)

### 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]
```

### 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.

### 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.

## RELATED DOCUMENTATION

| *Configuring PPPoE Service Name Tables*

## dynamic-profile (PPPoE Service Name Tables)

### 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]
```

### Release Information

Statement introduced in Junos OS Release 10.2.

### 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.

## 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

### Syntax

```
egress-policer-overhead bytes;
```

### Hierarchy Level

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

### Release Information

Statement introduced before Junos OS Release 11.1.

### 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.

## 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)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers

(**ethernet**, **vlan-ccc**, and **vlan-tcc** options only).

Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers. Only the **atm-ccc-cell-relay** and **atm-ccc-vc-mux** options are supported on ACX Series routers.

Statement introduced in Junos OS Release 17.3R1 for QFX10000 Series switches (**ethernet-ccc** and **vlan-ccc** options only).

### Description

Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

### 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.

## 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 | 49](#)

---

*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

### List of Syntax

[Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series on page 1150](#)

[Syntax for Physical Interfaces: SRX Series on page 1150](#)

[Syntax for Logical Interfaces: SRX Series on page 1150](#)

### 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 ( dtx | 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 ]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 11.1 for EX Series switches.




Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (**flexible-ethernet-services**, **ethernet-ccc**, and **ethernet-tcc** options only).

### 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

[Warning: element unresolved in stylesheets: <title> (in <config-options>). This is probably a new element that is not yet supported in the stylesheets.]

### 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.

## 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

### Syntax

```
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
      port-priority
    }
  }
  (auto-negotiation | no-auto-negotiation);
  ethernet-switch-profile {
    tag-protocol-id;
  }
  (flow-control | no-flow-control);
  ieee-802-3az-eee;
  link-mode mode;
  (loopback | no-loopback);
  speed (speed | auto-negotiation);
}
```

### Hierarchy Level

```
[edit interfaces interface-name],
[edit interfaces interface-range range]
```

### Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.3R2.

### Description

Configure Ethernet properties for a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface.

**NOTE:** The **ether-options** statement is not supported for subscriber management on aggregated Ethernet member link interfaces. You must configure **gether-options** instead.

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

---

*Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

---

*Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support*

---

*Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*

---

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

---

*Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support*

---

[Junos OS Ethernet Interfaces Configuration Guide](#)



## ethernet (Chassis)

### Syntax

```
ethernet {  
  device-count number;  
  lacp {  
    link-protection {  
      non-revertive;  
    }  
    system-priority;  
  }  
}
```

### Hierarchy Level

[edit chassis [aggregated-devices](#)]

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

### 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.

## RELATED DOCUMENTATION

[Configuring Junos OS for Supporting Aggregated Devices | 77](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## ethernet (Protocols OAM)

### List of Syntax

Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches on page 1160

Syntax: EX Series Switches and NFX Series Devices on page 1164

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

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



```

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;
  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;
  }
}

```

```
lmi {  
    status-counter count;  
    polling-verification-timer value;  
    interface name {  
        uni-id uni-name;  
        status-counter number;  
        polling-verification-timer value;  
        evc-map-type (all-to-one-bundling | bundling | service-multiplexing);  
        evc evc-name {  
            default-evc;  
            vlan-list vlan-id-list;  
        }  
    }  
}
```

## Syntax: EX Series Switches and NFX Series Devices

```

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

```

```

        priority priority-value;
    }
}
}
short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
}
}
performance-monitoring {
    sla-iterator-profiles {
        profile-name {
            calculation-weight {
                delay delay-value;
                delay-variation delay-variation-value;
            }
            cycle-time cycle-time-value;
            iteration-period iteration-period-value;
            measurement-type two-way-delay;
            passive;
        }
    }
}
}
traceoptions {
    file filename <files number> <match regex> <size size> <world-readable | no-world-readable>;
    flag flag ;
    no-remote-trace;
}
}

```

```

link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
  interface interface-name {
    link-discovery (active | passive);
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    negotiation-options {
      allow-remote-loopback;
      no-allow-link-events;
    }
  }
  traceoptions {
    file filename <files number> <match regex> <size size> <world-readable | no-world-readable>;
    flag flag ;
    no-remote-trace;
  }
}

```

## Hierarchy Level

[edit protocols [oam](#)]



**Release Information**

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

Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

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

**Description**

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

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

**Required Privilege Level**

interface—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

**RELATED DOCUMENTATION**

---

*Enabling IEEE 802.3ah OAM Support*

*Example: Configuring Ethernet OAM Link Fault Management*

## ethernet-policer-profile

### 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* gigether-options ethernet-switch-profile],  
 [edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

**RELATED DOCUMENTATION**

| [Configuring Gigabit Ethernet Policers](#) | 293

## evcs

### Syntax

```
evcs evc-id {  
    evc-protocol cfm;  
    remote-uni-count count;  
    multipoint-to-multipoint;  
}
```

### Hierarchy Level

[edit protocols [oam ethernet](#)]

### Release Information

Statement introduced in Junos OS Release 9.5.

### 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.

## RELATED DOCUMENTATION

[Configuring Ethernet Local Management Interface](#) | 579

[Imi \(Ethernet OAM\)](#) | 1241

## evc-protocol cfm

### Syntax

```
evc-protocol cfm {
  maintenance-association association-name | vpls routing-instance routing-id;
  maintenance-domain domain-id;
  mep-id mep-id;
}
```

### Hierarchy Level

[edit protocols [oam](#) ethernet [evcs](#)]

### Release Information

Statement introduced in Junos OS Release 9.5.

**mep-id** *mep-id* statement introduced in Junos OS Release 15.1.

### Description

Specify connectivity fault management (CFM) or virtual private LAN service (VPLS) as the Ethernet virtual connection (EVC) protocol.

### Options

**management-domain** *domain-id*—(Optional) For CFM, specify the CFM management domain.

**management-association** *association-id*—(Optional) For CFM, specify the CFM management association.

**routing-instance** *instance-id*—(Optional) For VPLS, specify the VPLS routing instance.

**mep-id** *mep-id*—(Required for CFM) Identifier for the maintenance association endpoint

**NOTE:** This option is available on MX Series routers only.

**Range:** 1 through 8191

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

Configuring Ethernet Local Management Interface | 579

Imi (Ethernet OAM) | 1241

## event (LFM)

### List of Syntax

[Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches on page 1173](#)

[Syntax: EX Series Switches and NFX Series Devices on page 1173](#)

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

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

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

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

### Hierarchy Level

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

### Release Information

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

Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX devices.

### Description

Configure link events in an action profile for Ethernet OAM link fault management (LFM).

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

**Required Privilege Level**

interface—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

**RELATED DOCUMENTATION**

[Monitoring Protocol Status | 676](#)

*Configuring Ethernet OAM Link Fault Management*



## event-thresholds

### Syntax

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

### Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### Description

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

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

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

| [Configuring Threshold Values for Local Fault Events on an Interface](#) | 662

## family

### 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;
  master-only;
  multipoint-destination address dlcid dlcid-identifier;
  multipoint-destination address {
    epd-threshold cells;
    inverse-arp;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Option **max-sessions-vs-a-ignore** introduced in Junos OS Release 11.4.

### Description

Configure protocol family information for the logical interface.

**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.

#### RELATED DOCUMENTATION

| *Configuring the Protocol Family*

## fastether-options

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.



RELATED DOCUMENTATION

| [Ethernet Interfaces Overview](#) | 3

## flow-control

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 in EX Series switches.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

### 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.

RELATED DOCUMENTATION

|                                                                                        |
|----------------------------------------------------------------------------------------|
| <a href="#">Configuring Flow Control   19</a>                                          |
| <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>                         |
| <i>Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support</i> |

## fnp

### Syntax

```
fnp {
  interval <100ms | 1s | 10s | 1m | 10m>;
  loss-threshold number
  interface interface name {
    domain-id domain-id
  }
}
```

### Hierarchy Level

[edit protocols [oam ethernet](#)]

### Release Information

Command introduced in Junos OS Release 11.4.

### 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.

### RELATED DOCUMENTATION

[Ethernet Failure Notification Protocol Overview](#) | 701

*Configuring the Failure Notification Protocol*

## force-up

### Syntax

```
force-up;
```

### Hierarchy Level

```
[edit interfaces interface-name aggregated-ethernet-options lacp]
```

### Release Information

Statement introduced in Junos OS Release 14.2 for MX Series routers.

**NOTE:** For EX9200 switches, you must configure **force-up** on physical interfaces of both MC-LAG peers for this feature to work properly.

### 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.

### RELATED DOCUMENTATION

*Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up*

## forwarding-class (Gigabit Ethernet IQ Classifier)

### Syntax

```
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 premium]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 293](#)

[input-priority-map | 1207](#)

**forwarding-class** statement in the *Class of Service User Guide (Routers and EX9200 Switches)*

## forwarding-mode (100-Gigabit Ethernet)

### 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]
```

### Release Information

Statement introduced in Junos OS Release 10.4.

Statement introduced in Junos OS Release 12.1 for MX Series routers.

### 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.

## RELATED DOCUMENTATION

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 256](#)

*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 | 263](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 264](#)

[sa-multicast \(100-Gigabit Ethernet\) | 1344](#)

[vlan-rule \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1424](#)

[vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1425](#)

## forwarding-mode (PTX Series Packet Transport Routers)

### Syntax

```
forwarding-mode {  
    sa-multicast  
}
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 12.1X48R4.

### 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.

## RELATED DOCUMENTATION

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 267](#)

[Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | 263](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 264](#)



## frame-error

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Threshold Values for Local Fault Events on an Interface | 662](#)

[Configuring Threshold Values for Fault Events in an Action Profile | 677](#)

## frame-period

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Threshold Values for Local Fault Events on an Interface](#) | 662

[Configuring Threshold Values for Fault Events in an Action Profile](#) | 677

## frame-period-summary

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Threshold Values for Local Fault Events on an Interface](#) | 662

[Configuring Threshold Values for Fault Events in an Action Profile](#) | 677

## framing (10-Gigabit Ethernet Interfaces)

### 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)
```

### Release Information

Statement introduced in Junos OS Release 8.0.

Statement introduced in Junos OS Release 12.3R2 for PTX Series Packet Transport Routers.

The option **precise-bandwidth** introduced in Junos OS Release 19.3R1 for 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.

**RELATED DOCUMENTATION**

[Framing Overview](#) | 225

*Configuring SONET Options for 10-Gigabit Ethernet Interfaces*

## gigether-options

### 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
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### Description

Configure Gigabit 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.

### RELATED DOCUMENTATION

[Ethernet Interfaces Overview](#) | 3

*gigether-options (ACX Series)*

## hold-time up

### Syntax

```
hold-time up timer-value;
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options lacp],
```

### Release Information

Statement introduced in Junos OS Release 14.2R3.

### 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.

## RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#) | 67

*Configuring Aggregated Ethernet LACP (CLI Procedure)*



## ieee802.1p

### Syntax

```
ieee802.1p premium [ values ];
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-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]
```

### Release Information

Statement introduced before Junos Release 7.4.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

### 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.

## RELATED DOCUMENTATION

| krupac

## ignore-l3-incompletes

### Syntax

```
ignore-l3-incompletes;
```

### Hierarchy Level

```
[edit interfaces interface-name fastether-options],  
[edit interfaces interface-name gigether-options]
```

### Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

### 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.

## RELATED DOCUMENTATION

| *Ignoring Layer 3 Incomplete Errors*

## ingress-policer-overhead

### Syntax

```
ingress-policer-overhead bytes;
```

### Hierarchy Level

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

### Release Information

Statement introduced before Junos OS Release 11.1.

Statement introduced in Junos OS Release 15.1X49-D30 for vSRX.

### 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.

## 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

### Syntax

```
ingress-rate-limit rate;
```

### Hierarchy Level

```
[edit interfaces interface-name fastether-options]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring the Ingress Rate Limit | 8](#)

## inline

### Syntax

```
inline;
```

### Hierarchy Level

```
[edit protocols lacp ppm]
```

### Release Information

Statement introduced in Junos OS Release 19.1R1 for MX Series routers with MPC line cards.

### 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 LCAP 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.

## RELATED DOCUMENTATION

[centralized](#) | [1127](#)

---

*Disabling or Enabling Inline Periodic Packet Management for LACP Packets*

---

*Configuring Link Aggregation*

---

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

## input-policer

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

### 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.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 293](#)

[input-three-color | 1208](#)

[layer2-policer | 1224](#)

*logical-interface-policer*

[output-policer | 1301](#)

[output-three-color | 1303](#)



## input-priority-map

### Syntax

```
input-priority-map {  
    ieee802.1p premium [ values ];  
}
```

### Hierarchy Level

```
[edit interfaces interface-name gige-ether-options ethernet-switch-profile ethernet-policer-profile]  
[edit interfaces interface-name ether-options ethernet-switch-profile ethernet-policer-profile]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

### 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.

### RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

[output-priority-map](#) | 1302

## input-three-color

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

### 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.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 293](#)

[input-policer | 1206](#)

[layer2-policer | 1224](#)

*logical-interface-policer*

[output-policer | 1301](#)

[output-three-color | 1303](#)

## input-vlan-map (Aggregated Ethernet)

### 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]
```

### 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.

### 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.

## RELATED DOCUMENTATION

*Stacking a VLAN Tag*

[output-vlan-map \(Aggregated Ethernet\)](#) | 1304

## interface

### Syntax

```
interface (all | interface-name) {  
    disable;  
}
```

### Hierarchy Level

```
[edit protocols lldp],  
[edit routing-instances routing-instance-name protocols lldp]
```

### Release Information

Statement introduced in Junos OS Release 9.6.

### Description

Specify an LLDP interface.

### Options

***interface-name***—A valid physical interface name.

**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, you can also specify LLDP neighbor details for management interfaces, such as **fxp** or **me**, on MX Series routers.

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*.

**all**—Run LLDP on all interfaces.

**disable**—Disable LLDP on the specified interface

### Required Privilege Level

routing—To view this statement in the configuration.

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

#### RELATED DOCUMENTATION

| *Configuring LLDP*

## interface (OAM Link-Fault Management)

### 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**]

### Release Information

Statement introduced in Junos OS Release 8.2.

### 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.

## RELATED DOCUMENTATION

| *Enabling IEEE 802.3ah OAM Support*

## interface-group

### 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 ]
```

### Release Information

Statement introduced in Junos OS Release 18.1R1.

### 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.

## RELATED DOCUMENTATION

[Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces](#) | 592

[interface-group-down](#) | 1215



## interface-group-down

### Syntax

```
interface-group-down
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management action-profile action-profile-name action]
```

### Release Information

Statement introduced in Junos OS Release 18.1R1.

### 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.

## RELATED DOCUMENTATION

[Configuring a CFM Action Profile to Bring Down a Group of Logical Interfaces | 592](#)

[interface-group | 1214](#)

## interface-none

### 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)

### 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 ],
```

### Release Information

Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers.

### 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.

## RELATED DOCUMENTATION

## lacp (802.3ad)

### Syntax

```
lacp {  
    port-priority port-priority;  
}
```

### Hierarchy Level

```
[edit interfaces interface-name fastether-options 802.3ad],  
[edit interfaces interface-name gigether-options 802.3ad]
```

### Release Information

Statement introduced in Junos OS Release 9.3.

### 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.

### RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#) | 67

[port-priority](#) | 1317

## lACP (Aggregated Ethernet)

### List of Syntax

[Syntax \(NFX Series\) on page 1219](#)

[Syntax \(EX Series\) on page 1219](#)

### Syntax (NFX Series)

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

### Syntax (EX Series)

```
lACP {
    (active | passive);
    admin-key key;
    accept-data;
    fast-failover;
    link-protection {
        disable;
        (revertive | non-revertive);
    }
    periodic interval;
    system-id mac-address;
    system-priority priority;
}
```

### Hierarchy Level (EX Series)

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

### Hierarchy Level (NFX Series)

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

### Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

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

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

### Description

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

For EX Series, when you configure the **accept-data** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

**NOTE:** When you configure the **accept-data** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

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

### Default

If you do not specify LACP as either **active** or **passive**, LACP remains passive.

### Options

**active**—Initiate transmission of LACP packets.

**admin-key *number***—Specify an administrative key for the router or switch.

**NOTE:** You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

**fast-failover**—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

**passive**—Respond to LACP packets.

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

### Required Privilege Level

**interface**—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

---

*Configuring Link Aggregation*

---

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

---

*Understanding Aggregated Ethernet Interfaces and LACP for Switches*

---

[Configuring LACP for Aggregated Ethernet Interfaces](#) | 67

## lACP (Protocols)

### 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]

### 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.

### 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.

#### RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#) | 67

*Unified ISSU System Requirements*

## layer2-policer

### 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],
```

### Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

### 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.

**RELATED DOCUMENTATION**

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Two-Color and Tricolor Policers](#) | 300

## link-adjacency-loss

**Syntax**

```
link-adjacency-loss;
```

**Hierarchy Level**

```
[edit protocols oam ethernet link-fault-management action-profile event]
```

**Release Information**

Statement introduced in Junos OS Release 8.5.

**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.

**RELATED DOCUMENTATION**

[Monitoring the Loss of Link Adjacency](#) | 675

## link-discovery

### Syntax

```
link-discovery (active | passive);
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

### Release Information

Statement introduced in Junos OS Release 8.2.

### 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.

### RELATED DOCUMENTATION

| [Configuring Link Discovery](#) | 661

## link-degrade-monitor

### Syntax

```
link-degrade-monitor {
  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]
```

### Release Information

Statement introduced in Junos OS Release 15.1.

### Description

Configure link degrade monitoring on an interface and specify the corrective action to be triggered when a link degrade event is detected. Deleting the configuration disables the feature. When configured, the feature monitors the quality of physical links on Ethernet interfaces (10-Gigabit, 40-Gigabit, and 100-Gigabit) and triggers the user-configured action when the link's bit error rate (BER) value breaches the preconfigured threshold. This feature can detect a BER value as low as  $10^{-13}$  to  $10^{-5}$ .

### Options

**actions media based**—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.

## RELATED DOCUMENTATION

[Link Degrad Monitoring Overview | 843](#)

[thresholds | 1384](#)

[recovery | 1335](#)

[request interface link-degrade-recover | 1476](#)

## link-down

### Syntax

```
link-down;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management ]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### 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.

## RELATED DOCUMENTATION

[Specifying the Actions to Be Taken for Link-Fault Management Events | 674](#)

## link-event-rate

### Syntax

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

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 8.5.

### 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.

### RELATED DOCUMENTATION

| [Configuring Threshold Values for Fault Events in an Action Profile](#) | 677

## link-fault-management

### 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](#)]

**Release Information**

Statement introduced in Junos OS Release 8.2.

**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.

**RELATED DOCUMENTATION**

| *Enabling IEEE 802.3ah OAM Support*

## link-mode

### 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]
```

### 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 12.2 for ACX Series Universal Metro Routers.

### 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.

**RELATED DOCUMENTATION**

[Configuring the Link Characteristics on Ethernet Interfaces | 9](#)

---

*Understanding Management Ethernet Interfaces*

---

*Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

---

*Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support*

## link-protection

### Syntax

```
link-protection {
  disable;
  (revertive | non-revertive);
}
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
[edit interfaces aex aggregated-ether-options lacp]
```

### Release Information

Statement introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

Support for **disable**, **revertive**, and **non-revertive** statements added in Junos OS Release 9.3.

### 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 gige-ether-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.

## RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Link Protection | 133](#)

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## link-protection (non-LACP)

### Syntax

```
link-protection {
    link-protection-revertive;
}
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

### Release Information

Statement introduced in Junos OS Release 17.3R1.

### 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 gige-ether-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.

RELATED DOCUMENTATION

| [Configuring Aggregated Ethernet Link Protection](#) | 133

## link-speed (Aggregated Ethernet)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

**mixed** option added in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers and 15.1F6 and 16.1R2 for PTX3000 routers.

### 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**—Links are of various speeds.
- **oc192**—Links are OC192.

**mixed**—Enables bundling of different Ethernet rate links in the same Aggregated Ethernet interface.

#### Required Privilege Level

interface—To view this statement in the configuration.

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

#### RELATED DOCUMENTATION

---

[Aggregated Ethernet Interfaces Overview | 58](#)

---

[Configuring Aggregated Ethernet Link Speed | 81](#)

---

[Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles | 94](#)

---

*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)

### Syntax

```
link-speed (speed | mixed);
```

### Hierarchy Level

```
[edit interfaces asx aggregated-sonet-options]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

**mixed** option added in Release 8.0.

### 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.

## RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Link Speed](#) | 81

*Configuring Aggregated SONET/SDH Interfaces*

## Imi (Ethernet OAM)

### Syntax

```

Imi {
  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](#)]

### Release Information

Statement introduced in Junos OS Release 9.5.

### 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.

#### RELATED DOCUMENTATION

[Configuring Ethernet Local Management Interface | 579](#)

[evcs | 1170](#)

## load-balance

### 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]
```

### Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Release 14.1 for PTX Series Packet Transport Routers.

### 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.

**RELATED DOCUMENTATION**

[Understanding Aggregated Ethernet Load Balancing | 141](#)

[Example: Configuring Aggregated Ethernet Load Balancing | 161](#)

## load-balance-stateful (Aggregated Ethernet Interfaces)

### Syntax

```
load-balance-stateful {
  per-flow;
  rebalance interval;
  load-type (low | medium | large);
}
```

### Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options]
```

### Release Information

Statement introduced in Junos OS Release 13.2R1.

### 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.

### RELATED DOCUMENTATION

[Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces](#) | 147

## load-type (Aggregated Ethernet Interfaces)

### Syntax

```
load-type (low | medium | large);
```

### Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful]
```

### Release Information

Statement introduced in Junos OS Release 13.2R1.

### 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.

## RELATED DOCUMENTATION

[Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces](#) | 147



## local-bias (ae load-balance)

### Syntax

```
local-bias percent bias;
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options load-balance]
```

### Release Information

Statement introduced in Junos OS Release 19.2R1.

### 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 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.

**RELATED DOCUMENTATION**

[Understanding Aggregated Ethernet Load Balancing | 141](#)

[Example: Configuring Aggregated Ethernet Load Balancing | 161](#)

## logical-tunnel-options

### 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]
```

### Release Information

Statement introduced in Junos OS Release 19.2R1.

### 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

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

### 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 gigether-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]
```

### Release Information

Statement introduced before Junos OS Release 7.4 for MX Series.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

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

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

Statement modified in Junos OS Release 9.2 for the SRX Series.

### 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 gigether-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.

**RELATED DOCUMENTATION**

*Configuring Ethernet Loopback Capability*

*Understanding Interfaces*

## loopback (Local and Remote)

### Syntax

```
loopback (local | remote);
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

### Release Information

Statement introduced in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers.

Statement introduced in Junos OS Release 15.1F6 and 16.1R2 for PTX3000 routers.

### 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

**local**—Enables local loopback

**remote**—Enables remote loopback

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

*Configuring Ethernet Loopback Capability*

## loopback-tracking

### Syntax

```
loopback-tracking;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### 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.

## RELATED DOCUMENTATION

[IEEE 802.3ah OAM Link-Fault Management Overview](#) | 655

*Enabling IEEE 802.3ah OAM Support*



## loss-priority

### Syntax

```
loss-priority (high | low);
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map  
  classifier premium forwarding-class class-name]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

## mac

### Syntax

```
mac mac-address;
```

### Hierarchy Level

```
[edit interfaces interface-name]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### Description

Set the MAC address of the interface.

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

### Options

**mac-address**—MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*. For example, **0000.5e00.5355** or **00:00:5e:00:53:55**.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring the MAC Address on the Management Ethernet Interface | 28](#)

*Configuring a Pseudowire Subscriber Logical Interface Device*

## mac-address (Accept Source Mac)

### 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 ]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

### 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 \(MAC\)](#).

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

## mac-learn-enable

### Syntax

```
mac-learn-enable;
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]  
[edit interfaces aex aggregated-ether-options ethernet-switch-profile]
```

### 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.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers | 293](#)

[Configuring MAC Address Accounting | 24](#)

## mac-validate

### Syntax

```
mac-validate (loose | strict);
```

### Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family]
```

### Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

### 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.

## RELATED DOCUMENTATION

[MAC Address Validation on Static Ethernet Interfaces Overview](#) | 11

[Configuring an IP Demultiplexing Interface](#)

[Configuring a VLAN Demultiplexing Interface](#)

## master-only

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

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

### Description

Configure the IP address to be used when the Routing Engine is the current master.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring a Consistent Management IP Address | 26](#)

*CLI User Guide*

## max-sessions (PPPoE Service Name Tables)

### Syntax

```
max-sessions number;
```

### Hierarchy Level

```
[edit protocols pppoe service-name-tables table-name service service-name]
```

### Release Information

Statement introduced in Junos OS Release 10.2.

### 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.

## 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)

### 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]
```

### Release Information

Statement introduced in Junos OS Release 11.4.

### 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-vsa-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.

## 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

### Syntax

```
maximum-links maximum-links-limit;
```

### Hierarchy Level

```
[edit chassis aggregated-devices]
```

### Release Information

Statement introduced in Junos OS Release 11.1 for T Series routers.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Statement introduced in Junos OS Release 12.3 for MX Series routers.

### 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.

### RELATED DOCUMENTATION

---

*Network Services Mode Overview*

---

[Configuring Junos OS for Supporting Aggregated Devices | 77](#)

---

[Configuring an Aggregated Ethernet Interface | 64](#)

---

*network-services*

## mc-ae

### 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]
```

### Release Information

Statement introduced in Junos OS Release 9.6 for MX Series routers.

**events** statement introduced in Junos OS Release 11.4R4 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series. Only the **chassis-id**, **mc-ae-id**, **mode active-active**, and **status-control (active | standby)** options are supported on QFX Series devices.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**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.

### 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

**chassis-id**—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.

**Values:** 0 or 1

**events**—Specify an action if a specific MC-LAG event occurs.

**iccp-peer-down**—Specify an action if the ICCP peer of this node goes down.

**force-icl-down**—If the node's ICCP peer goes down, bring down the interchassis-link logical interface.

**prefer-status-control-active**—Specify that the node configured as **status-control active** become the active node if the peer of this node goes down.

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.

**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.

**Range:** 1 through 4,294,967,294



**revert-time**—Wait interval (in minutes) before the switchover to the preferred node is performed when the **switchover-mode** is configured as revertive.

**Range:** 1 through 10

**status-control (active | standby)**—Specify whether the chassis becomes active or remains in standby mode when an interchassis link failure occurs.

- Events ICCP-Peer-Down Force-ICL-Down

Forces the ICL to be down if the peer of this node goes down.

- Events ICCP-Peer-Down Prefer-Status-Control-Active

Allows the LACP system ID to be retained during a reboot, which provides better convergence after a failover.

**switchover-mode (non-revertive | revertive)**—Specify whether Junos OS should trigger a link switchover to the preferred node when the active node is available.

**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.

## minimum-bandwidth (aggregated Ethernet)

### Syntax

```
minimum-bandwidth bw-unit unit bw-value value;
```

### Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

### Release Information

Statement introduced before Junos OS Release 14.1R1 and 14.2 for MX Series.

### 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.

### 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.

### Required Privilege Level

routing—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Aggregated Ethernet Interfaces Overview | 58](#)

[Understanding Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles | 86](#)

[Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles | 94](#)

## minimum-links

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]
```

### 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.

### 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.

### 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.

**RELATED DOCUMENTATION**

[Configuring Aggregated Ethernet Minimum Links | 135](#)

[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****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)
```

**Release Information**

Statement introduced in Junos OS Release 13.3.

**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.

**RELATED DOCUMENTATION**

[Modes of Operation | 228](#)

[Configuring Mixed-Rate Mode Operation | 229](#)

## mtu

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

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.

Statement introduced in Junos OS 17.3R1 Release for MX Series Routers.

### 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  $[6(\text{DMAC})+6(\text{SMAC})+2(\text{EtherType})]$ . 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(DMAC)+6(SMAC)+2(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 (**SVLAN+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.

**RELATED DOCUMENTATION**

---

*Configuring the Media MTU*

---

*Configuring the MTU for Layer 2 Interfaces*

---

*Setting the Protocol MTU*



## mru

### 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]
```

### Release Information

Statement introduced in Junos OS Release 19.1R1 for MX Series Routers.

### 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

*mru*—MRU size in bytes.

**Range:** 256 through 16008 bytes.

### Required Privilege Level

interface

### RELATED DOCUMENTATION

[Ethernet Interfaces Overview](#) | 3

[gigether-options](#) | 1196

## multicast-statistics

### Syntax

```
multicast-statistics;
```

### Hierarchy Level

```
[edit interfaces interface-name]
```

### Release Information

Statement introduced before Junos OS Release 10.2.

### 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.

## RELATED DOCUMENTATION

---

[Configuring Multicast Statistics Collection on Aggregated Ethernet Interfaces | 84](#)

---

*Configuring Multicast Statistics Collection on Aggregated SONET Interfaces*

---

[Configuring Multicast Statistics Collection on Ethernet Interfaces | 10](#)

---

*Configuring Multicast Statistics Collection on SONET Interfaces*

## negotiate-address

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

*Configuring IPCP Options for Interfaces with PPP Encapsulation*

[address](#) | [1097](#)

[unnumbered-address \(PPP\)](#) | [1420](#)

*Junos OS Administration Library*

## negotiation-options

### Syntax

```
negotiation-options {  
    allow-remote-loopback;  
    no-allow-link-events;  
}
```

### Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

### RELATED DOCUMENTATION

| [IEEE 802.3ah OAM Link-Fault Management Overview](#) | 655

## no-adaptive

### 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]
```

### Release Information

Statement introduced in Junos OS Release 13.2R3.

### 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.

### RELATED DOCUMENTATION

[Understanding Aggregated Ethernet Load Balancing](#) | 141

## no-allow-link-events

### Syntax

```
no-allow-link-events;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name negotiation-options]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

### RELATED DOCUMENTATION

[Disabling the Sending of Link Event TLVs](#) | 663

## no-auto-mdix

### Syntax

```
no-auto-mdix;
```

### Hierarchy Level

```
[edit interface ge-fpc/port/pic together-options]
```

### Release Information

Statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

### 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.

## RELATED DOCUMENTATION

[Ethernet Interfaces Overview](#) | 3

[gether-options](#) | 1196.

## no-keepalives

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

---

*Configuring Keepalives*

---

[Disabling the Sending of PPPoE Keepalive Messages](#) | 52

---

*Configuring Frame Relay Keepalives*



## no-pre-classifier

### Syntax

```
no-pre-classifier;
```

### Hierarchy Level

```
[edit chassis fpc n pic n]
```

### Release Information

Statement introduced in Junos OS Release 10.4.

### 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.

## RELATED DOCUMENTATION

[10-port 10-Gigabit Ethernet LAN/WAN PIC Overview | 200](#)

[Disabling Control Queue Disable on a 10-port 10-Gigabit Ethernet LAN/WAN PIC | 232](#)

## no-send-pads-ac-info

### Syntax

```
no-send-pads-ac-info;
```

### Hierarchy Level

```
[edit protocols pppoe]
```

### Release Information

Statement introduced in Junos OS Release 12.2.

### 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.

### RELATED DOCUMENTATION

*Disabling the Sending of PPPoE Access Concentrator Tags in PADS Packets*

## no-send-pads-error

### Syntax

```
no-send-pads-error;
```

### Hierarchy Level

```
[edit protocols pppoe]
```

### Release Information

Statement introduced in Junos OS Release 12.3.

### 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.

## RELATED DOCUMENTATION

*Discarding PADR Messages to Accommodate Abnormal CPE Behavior*

## non-revertive (Interfaces)

### Syntax

```
non-revertive;
```

### Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

### Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### 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.

## RELATED DOCUMENTATION

[link-protection](#) | 1234

[Configuring Aggregated Ethernet Link Protection](#) | 133

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## number-of-sub-ports

### Syntax

```
number-of-sub-ports <number-of-sub-ports>;
```

### Hierarchy Level

```
[edit chassis fpc fpc-slot pic pic-number port port-num]
```

### Release Information

Statement introduced in Junos OS Release 19.1R1 for MPC10E-15C-MRATE supported on MX240, MX480, and MX960 routers.

Statement introduced in Junos OS Evolved Release 19.1R1 for PTX10003-80C and PTX10003-160C router.

### 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.

(MPC11E) To specify the number of interfaces to be created on a physical port.

### Options

On PTX10003-80C, PTX10003-160C router: *number-of-sub-ports*—Specify the number of sub-ports per physical port. Configurable options are 1, 2, and 4.

On MPC10E-15C-MRATE supported on MX240, MX480, and MX960 routers:

*number-of-sub-ports*—Specify the number of sub-ports per physical port. Configurable options are 1, 2, 3, and 4.

**NOTE:** You can configure the *number-of-sub-ports* only for 10-Gbps speed. For other speeds, this configuration is not supported.

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[speed \(Ethernet\) | 1359](#)

[PTX10003 Router Rate-Selectability Overview | 340](#)

## oam

### 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;
    }
    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
    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]

**Release Information**

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

**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.

**RELATED DOCUMENTATION**

| [IEEE 802.3ah OAM Link-Fault Management Overview](#) | 655

## optics-options

### Syntax

```
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

**alarm** option and **warning** options introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

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.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 18.3R1 for ACX6360 routers.

**loopback** option introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

### 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

**loopback**—Displays the electrical loopback status of QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

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

**Required Privilege Level**

interface—To view this statement in the configuration.

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

**RELATED DOCUMENTATION**

---

[Ethernet DWDM Interface Wavelength Overview | 505](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[Supported Forward Error Correction Modes on ACX6360 Router | 504](#)

## otn-options

### 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]
```

## 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.

Statement introduced in Junos OS Release 19.2R1-S1 for ACX5448-D routers.

## 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.

## RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview](#) | 424

[100-Gigabit Ethernet OTN Options Configuration Overview](#) | 424

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#) | 512

## output-policer

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

### 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.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 293](#)

[input-policer | 1206](#)

[input-three-color | 1208](#)

[layer2-policer | 1224](#)

*logical-interface-policer*

[output-three-color | 1303](#)

## output-priority-map

### Syntax

```
output-priority-map {
  classifier {
    premium {
      forwarding-class class-name {
        loss-priority (high | low);
      }
    }
  }
}
```

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

### 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.

### RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

[input-priority-map](#) | 1207



## output-three-color

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

### 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.

## RELATED DOCUMENTATION

*Two-Color and Three-Color Policers at Layer 2*

*Applying Layer 2 Policers to Gigabit Ethernet Interfaces*

[Configuring Gigabit Ethernet Policers | 293](#)

[input-three-color | 1208](#)

[input-policer | 1206](#)

[layer2-policer | 1224](#)

*logical-interface-policer*

[output-policer | 1301](#)

## output-vlan-map (Aggregated Ethernet)

### 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]
```

### 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.

### 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.

### RELATED DOCUMENTATION

*Stacking and Rewriting Gigabit Ethernet VLAN Tags*

[input-vlan-map \(Aggregated Ethernet\)](#) | 1209

## pado-advertise

### Syntax

```
pado-advertise;
```

### Hierarchy Level

```
[edit protocols pppoe]
```

### Release Information

Statement introduced in Junos OS Release 10.2.

### 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.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

*Enabling Advertisement of Named Services in PADO Control Packets*

## passive-monitor-mode

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

*Enabling Passive Monitoring on ATM Interfaces*

[Passive Monitoring on Ethernet Interfaces Overview](#) | 839

*Enabling Packet Flow Monitoring on SONET/SDH Interfaces*

*multiservice-options*

*Junos OS Services Interfaces Library for Routing Devices*

## pdu-interval

### Syntax

```
pdu-interval interval;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

### Release Information

Statement introduced in Junos OS Release 8.2 for MX, M, T, ACX, Series routers, SRX Series firewalls, and EX Series Switches.

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

### Description

For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

### Options

**interval**—Periodic OAM PDU sending interval.

**Range:** For MX, M, T, ACX, Series routers, SRX Series firewalls and EX Series switches – 100 through 1000 milliseconds

**Default:** 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.

## RELATED DOCUMENTATION

[Configuring the OAM PDU Interval | 662](#)

*Example: Configuring Ethernet OAM Link Fault Management*

*Configuring Ethernet OAM Link Fault Management*

## pdu-threshold

### Syntax

```
pdu-threshold threshold-value;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

### Release Information

Statement introduced in Junos OS Release 8.2 for T, M, MX and ACX Series routers, SRX Series firewalls and EX Series switches.

Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

### Description

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.

## RELATED DOCUMENTATION

[Configuring the OAM PDU Threshold | 662](#)

[Configuring Ethernet OAM Link Fault Management](#)

## per-flow (Aggregated Ethernet Interfaces)

### Syntax

```
per-flow;
```

### Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful]
```

### Release Information

Statement introduced in Junos OS Release 13.2R1.

### 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.

### RELATED DOCUMENTATION

[Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces](#) | 147

## periodic

### List of Syntax

[Syntax \(EX Series\) on page 1310](#)

[Syntax \(QFX Series\) on page 1310](#)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

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

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

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### Description

For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

### Options

***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.

## RELATED DOCUMENTATION

---

[Configuring LACP for Aggregated Ethernet Interfaces | 67](#)

---

*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*

## **policer (CFM Firewall)**

### **Syntax**

```
policer cfm-policer {  
  if-exceeding {  
    bandwidth-limit 8k;  
    burst-size-limit 2k;  
  }  
  then discard;  
}
```

### **Hierarchy Level**

```
[edit firewall]
```

### **Release Information**

Statement introduced in Junos OS Release 10.0.

### **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.

## **RELATED DOCUMENTATION**

[Configuring Rate Limiting of Ethernet OAM Messages | 576](#)

[policer \(CFM Global\) | 1063](#)

[policer \(CFM Session\) | 1064](#)

## policer (CoS)

### 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* gigether-options ethernet-switch-profile [ethernet-policer-profile](#)]

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION



## policer (MAC)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

### 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.

**RELATED DOCUMENTATION**

| [Configuring Gigabit Ethernet Policers](#) | 293

## port-priority

### Syntax

```
port-priority priority;
```

### Hierarchy Level

```
[edit interfaces interface-name gigether-options 802.3ad lACP]
```

### Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### 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.

## RELATED DOCUMENTATION

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

*Configuring Aggregated Ethernet LACP (CLI Procedure)*

## pp0 (Dynamic PPPoE)

### 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 {
        input filter-name {
```



```

        precedence precedence;
    }
    output filter-name {
        precedence precedence;
    }
}
}
}
}
}

```

### Hierarchy Level

[edit dynamic-profiles *profile-name* interfaces]

### Release Information

Statement introduced in Junos OS Release 10.1.

### Description

Configure the dynamic PPPoE logical interface in a dynamic profile. When the router creates a dynamic PPPoE logical interface on an underlying Ethernet interface configured with PPPoE (**ppp-over-ether**) encapsulation, it uses the information in the dynamic profile to determine the properties of the dynamic PPPoE logical 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.

### Required Privilege Level

interface—To view this statement in the configuration.

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

### RELATED DOCUMENTATION

[Configuring a PPPoE Dynamic Profile](#)

[Configuring Dynamic Authentication for PPP Subscribers](#)

For information about creating static PPPoE interfaces, see [Configuring PPPoE | 45](#)

## ppm (Ethernet Switching)

### Syntax

```
ppm {  
    centralized;  
}
```

### Hierarchy Level

```
[edit protocols lacp]
```

### Release Information

Statement introduced in Junos OS Release 9.4 for MX Series routers.

Statement introduced in Junos OS Release 10.2 for EX Series switches.

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

Statement introduced in Junos OS Release 12.1 for T Series devices.

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

### 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.

## RELATED DOCUMENTATION

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

*Configuring Distributed Periodic Packet Management*

## pppoe-options

### 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],
```

### 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.

### 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.

#### RELATED DOCUMENTATION

| [Configuring a PPPoE Interface](#) | 45

## pppoe-underlying-options (Static and Dynamic Subscribers)

### Syntax

```
pppoe-underlying-options {
  access-concentrator name;
  dynamic-profile profile-name;
  direct-connect
  duplicate-protection;
  max-sessions number;
  max-sessions-vsa-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]
```

### Release Information

Statement introduced in Junos OS Release 10.0.

### 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.

## RELATED DOCUMENTATION

[Configuring PPPoE | 45](#) (for static interfaces)

[Configuring an Underlying Interface for Dynamic PPPoE Subscriber Interfaces](#)

[Assigning a Service Name Table to a PPPoE Underlying Interface](#)

## preferred-source-address

### 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],
```

### 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.

### 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.

**RELATED DOCUMENTATION**

---

*Configuring an Unnumbered Interface*

---

*Junos OS Network Interfaces Library for Routing Devices*

---

*Junos OS Administration Library*

## premium (Output Priority Map)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

[input-priority-map](#) | 1207



## premium (Policer)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

---

[Configuring Gigabit Ethernet Policers](#) | 293

---

[aggregate \(Gigabit Ethernet CoS Policer\)](#) | 1105

---

[ieee802.1p](#) | 1199

## protocol-down

### Syntax

```
protocol-down;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile event]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### 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.

### RELATED DOCUMENTATION

[Configuring an OAM Action Profile](#) | 673

## premium (Output Priority Map)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Policers](#) | 293

[input-priority-map](#) | 1207

## premium (Policer)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

---

[Configuring Gigabit Ethernet Policers](#) | 293

---

[aggregate \(Gigabit Ethernet CoS Policer\)](#) | 1105

---

[ieee802.1p](#) | 1199

# proxy

## 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]
```

## Release Information

Statement introduced before Junos OS Release 7.4.

## 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.

## 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.

## RELATED DOCUMENTATION

|                                                                  |
|------------------------------------------------------------------|
| <i>Configuring Translation Cross-Connect Interface Switching</i> |
| <i>Example: Configuring an Ethernet TCC or Extended VLAN TCC</i> |
| <i>remote</i>                                                    |
| <i>Junos OS VPNs Library for Routing Devices</i>                 |

## rebalance (Aggregated Ethernet Interfaces)

### Syntax

```
rebalance interval
```

### Hierarchy Level

```
[edit interfaces aeX unit logical-unit-number forwarding-options load-balance-stateful per-flow]
```

### Release Information

Statement introduced in Junos OS Release 13.2R1.

### 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.

### RELATED DOCUMENTATION

[Configuring Stateful Load Balancing on Aggregated Ethernet Interfaces](#) | 147

## receive-options-packets

### Syntax

```
receive-options-packets;
```

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

*Enabling Passive Monitoring on ATM Interfaces*

*Enabling Packet Flow Monitoring on SONET/SDH Interfaces*

## receive-ttl-exceeded

### Syntax

```
receive-ttl-exceeded;
```

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

*Enabling Passive Monitoring on ATM Interfaces*

*Enabling Packet Flow Monitoring on SONET/SDH Interfaces*



## recovery

### Syntax

```
recovery {  
  (auto | manual);  
  timer timer-value;  
}
```

### Hierarchy Level

```
[edit interfaces interfaces-name link-degrade-monitor]
```

### Release Information

Statement introduced in Junos OS Release 15.1.

### 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.

**RELATED DOCUMENTATION**

---

[Link Degrade Monitoring Overview | 843](#)

---

[link-degrade-monitor | 1227](#)

---

[thresholds | 1384](#)

---

[request interface link-degrade-recover | 1476](#)

## remote-loopback

### Syntax

```
remote-loopback;
```

### Hierarchy Level

```
[edit protocols oam link-fault-management interface interface-name]
```

### Release Information

Statement introduced in Junos OS Release 8.2.

### 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.

### RELATED DOCUMENTATION

| [Setting a Remote Interface into Loopback Mode](#) | 680

## restore-interval

### Syntax

```
restore-interval number;
```

### Hierarchy Level

```
[edit protocols protection-group ethernet-ring ring-name]
```

### 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.

### 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.

## 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

### Syntax

```
revertive;
```

### Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
```

### Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Release 12.3 for EX Series switches.

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### 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.

## RELATED DOCUMENTATION

*non-revertive (Chassis)*

*Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches*

## routing-instance

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

| *Junos OS Services Interfaces Library for Routing Devices*

## routing-instance (PPPoE Service Name Tables)

### 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]
```

### Release Information

Statement introduced in Junos OS Release 10.2.

### 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.

## 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

### Syntax

```
expected-defect {  
  rx-enable ;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

### Release Information

Statement introduced in Junos OS Release 19.1.

### 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.

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | **1132**

[show oam ethernet connectivity-fault-management mep-database](#) | **1901**



## rx-max-duration

### Syntax

```
expected-defect {
  rx-max-duration ;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

### Release Information

Statement introduced in Junos OS Release 19.1.

### 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*.

**Maximum 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.

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | **1132**

[show oam ethernet connectivity-fault-management mep-database](#) | **1901**

## sa-multicast (100-Gigabit Ethernet)

### Syntax

```
sa-multicast;
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 10.4.

### 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.

## RELATED DOCUMENTATION

[Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and PF-1CGE-CFP | 263](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs PF-1CGE-CFP and PD-1CE-CFP-FPC4 | 264](#)

*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 | 266](#)

---

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 267](#)

---

[forwarding-mode \(100-Gigabit Ethernet\) | 1189](#)

---

[sa-multicast \(PTX Series Packet Transport Routers\) | 1346](#)

---

[vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1425](#)

---

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 256](#)

---

## sa-multicast (PTX Series Packet Transport Routers)

### Syntax

```
sa-multicast;
```

### Hierarchy Level

```
[edit chassis fpc slot pic slot port port-number forwarding-mode]
```

### Release Information

Statement introduced in Junos OS Release 12.1X48R4.

### 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.

## RELATED DOCUMENTATION

[Interoperability Between the 100-Gigabit Ethernet PICs PD-1CE-CFP-FPC4 and P1-PTX-2-100GE-CFP | 266](#)

[Configuring the Interoperability Between the 100-Gigabit Ethernet PICs P1-PTX-2-100GE-CFP and PD-1CE-CFP-FPC4 | 267](#)

## send-critical-event

### Syntax

```
send-critical-event;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile action]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### 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.

### RELATED DOCUMENTATION

| [Specifying the Actions to Be Taken for Link-Fault Management Events](#) | 674

## server

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.5.

### 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.

## RELATED DOCUMENTATION

| [Configuring the PPPoE Server Mode](#) | 50

## service (PPPoE)

### 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]
```

### 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.

### 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.

#### 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

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

[Configuring the PPPoE Service Name](#) | 49

*Junos OS Interfaces and Routing Configuration Guide*

## service-name-table

### 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]
```

### Release Information

Statement introduced in Junos OS Release 10.0.

Support at the **[edit ... family pppoe]** hierarchies introduced in Junos OS Release 11.2.

### 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.

## 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

### 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]

### 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.

### 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.

#### RELATED DOCUMENTATION

[Configuring PPPoE Service Name Tables](#)

[Creating a Service Name Table](#)

## session-expiry (MX Series in Enhanced LAN Mode)

#### Syntax

```
session-expiry seconds;
```

#### Hierarchy Level

```
[edit protocols authentication-access-control interface (all | [interface-names])]
```

#### Release Information

Statement introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

#### 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.

## source-address-filter

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

### 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: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.

#### RELATED DOCUMENTATION

---

[Configuring MAC Address Filtering for Ethernet Interfaces | 21](#)

---

[Configuring MAC Address Filtering on PTX Series Packet Transport Routers | 23](#)

---

[source-filtering | 1358](#)

## source-filtering

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.

### 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.

## RELATED DOCUMENTATION

---

[Configuring MAC Address Filtering for Ethernet Interfaces](#) | 21

---

[Configuring MAC Address Filtering on PTX Series Packet Transport Routers](#) | 23

---

[accept-source-mac](#) | 1088

---

[source-address-filter](#) | 1356



## speed (Ethernet)

### List of Syntax

[Syntax \(EX Series\) on page 1359](#)

[Syntax \(EX2300 and EX4300\) on page 1359](#)

[Syntax \(EX Series, ACX Series, MX Series\) on page 1359](#)

[Syntax \(ACX5448\) on page 1359](#)

[Syntax \(ACX5448-D\) on page 1359](#)

[Syntax \(QFX Series, OCX1100, EX4600\) on page 1359](#)

[Syntax \(PTX10003-80C, PTX10003-160C\) on page 1359](#)

### Syntax (EX Series)

```
speed (auto-negotiation | speed) ;
```

### Syntax (EX2300 and EX4300)

```
speed speed;
```

### Syntax (EX Series, ACX Series, MX Series)

```
speed (10m | 10g | 100m | 1g | 2.5g | 5g | auto | auto-10m-100m);
```

### Syntax (ACX5448)

```
speed (100m | 1g | auto);
```

### Syntax (ACX5448-D)

```
speed (10g | 25g | 40g | 100g);
```

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

```
speed (10g | 1g | 100m)
```

### Syntax (PTX10003-80C, PTX10003-160C)

```
speed (10g | 25g | 40g | 100g | 400g)
```

### Hierarchy Level (EX Series)

```
[edit interfaces interface-name ether-options]
```

### Hierarchy Level (EX2300 and EX4300)

```
[edit interfaces interface-name]
```

### Hierarchy Level (ACX5448, ACX5448-D)

```
[edit interfaces interface-name]
```

### 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 (QFX Series, EX4600, OCX Series)

```
[edit interfaces interface-name]
```

### 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.

## Description

Configure the interface speed. This statement applies to the management Ethernet interface (**fxp0** or **em0**), Fast Ethernet 12-port and 48-port PICs, the built-in Fast Ethernet port on the FIC (M7i router), Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers, and Gigabit Ethernet interfaces on EX Series switches.

When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled. When you configure 100BASE-FX SFP, you must set the port speed at 100 Mbps.

**NOTE:** On MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the **auto-10m-100m** option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link. However, when you view the interface speed of the DPC, using the **show interfaces** command, the value of the speed is not accurately displayed. For instance, if you configure the speed of the Tri-rate enhanced DPC, as 100Mbps on both sides of the link, the interface speed of the DPC is negotiated to 100 Mbps. However, the interface speed of the DPC displays 1 bps. This is an issue with the **show interfaces** command only. The actual interface speed is 100 Mbps.

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

(For QFX5100-48T only) To negotiate any speed on QFX5100-48T switches:

```
set interfaces xe-0/0/0 ether-options auto-negotiation
```

```
set interfaces xe-0/0/0 speed auto
```

Basically, when you configure a port using **speed auto** option, the port deletes the last configured speed, comes up again and advertises all the possible speeds:

For a port to start with a specific speed, it is mandatory that both the auto-negotiation must be enabled and interface must be configured with a particular speed. Otherwise, the switch will remain with the last negotiated speed. If you only set the **auto-negotiation** option (and no speed option), then the switch will start with the last speed it connected at and will only advertise that to the server.

To configure a particular speed on QFX5100-48T switches, mention the speed:

```
set interfaces xe-0/0/0 ether-options auto-negotiation
```

```
set interfaces xe-0/0/0 speed speed
```

For example to configure 1-Gbps speed in full duplex mode, execute the following command:

```
set interfaces xe-0/0/0 ether-options auto-negotiation
```

```
set interfaces xe-0/0/0 speed 1g
```

Note that you can not disable auto-negotiation on 1-Gigabit Ethernet ports. It is mandatory to enable autonegotiation when 1-Gbps speed is configured on a particular interface.

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

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

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

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

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

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

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

(For ACX5448 only) 100-Gbps speed is supported from interfaces xe-0/0/24 to xe-0/0/47 only.

On ACX5448-D routers, you can configure speeds of 100-Gbps, 40-Gbps, 25-Gbps, and 10-Gbps are supported on the et- and ot- interfaces.

#### **Default (EX Series)**

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

**Options**

You can specify the speed as either **10m** (10 Mbps), **100m** (100 Mbps), and on MX Series routers, **1g** (1 Gbps). You can also specify the **auto** option on MX Series routers.

For Gigabit Ethernet interfaces on EX Series switches, you can specify one of the following options:

Table 121: Options for speed

| Platforms                           | Speed Supported                                                                                                                                                                                                                           | Auto-negotiation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EX Series Switches                  | <b>100m</b> —100 Mbps<br><b>10m</b> —10 Mbps<br><b>1g</b> —1 Gbps                                                                                                                                                                         | auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the <a href="#">auto-negotiation</a> statement at the <b>[edit interfaces interface-name ether-options]</b> hierarchy level is enabled.                                                                                                                                                                                                                                                               |
| ACX, MX Series                      | <b>100m</b> —100 Mbps<br><b>10m</b> —10 Mbps<br><b>1g</b> —1 Gbps                                                                                                                                                                         | <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.<br><br><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.                                                                                                                                                                                                                                                                         |
| EX4600, QFX Series, QFabric, OCX100 | <b>10g</b> —10 Gbps<br><b>1g</b> —1 Gbps<br><b>100m</b> —100 Mbps                                                                                                                                                                         | auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the <a href="#">auto-negotiation</a> statement at the <b>[edit interfaces interface-name ether-options]</b> hierarchy level is enabled.                                                                                                                                                                                                                                                               |
| PTX10003-80C, and PTX10003-160C     | <b>10g</b> —10 Gbps<br><b>40g</b> —40 Gbps<br><b>100g</b> —100 Gbps                                                                                                                                                                       | <i>speed</i> —Specify the interface speed.<br><br>Use <b>number-of-sub-ports</b> configuration statement 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.<br><br>See <a href="#">“PTX10003 Router Rate-Selectability Overview” on page 340</a> for more details.                                                                  |
| EX2300                              | <b>10m</b> —10 Mbps (supported on EX series switches and only on <b>ge</b> interfaces of EX2300 switch)<br><b>100m</b> —100 Mbps<br><b>1g</b> —1 Gbps<br><b>2.5g</b> —2.5 Gbps (supported only on <b>mge</b> interfaces of EX2300 switch) | <i>speed</i> —Specify the interface speed. If the <a href="#">auto-negotiation</a> statement at the <b>[edit interfaces interface-name ether-options]</b> hierarchy level is disabled, you must specify a specific value. This value sets the speed that is used on the link. If the <b>auto-negotiation</b> statement is enabled, you might want to configure a specific speed value to advertise the desired speed to the remote end.<br><br>The Multi-rate gigabit ethernet interface (MGE) on EX2300-24MP and EX2300-48MP switches flaps |

Table 121: Options for speed (*continued*)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | (becomes unavailable, and then available again) while performing timeout detection and recovery (TDR) test.                                                                                                                                                       |
| EX4300-48MP<br>(EX-UM-4SFPP-MR) | <p><b>10m</b>—10 Mbps (supported only on <b>ge</b> interfaces)</p> <p><b>100m</b>—100 Mbps (supported on <b>ge</b> and <b>mge</b> interfaces)</p> <p><b>1g</b>—1 Gbps (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.</p> <p><b>2.5g</b>—2.5 Gbps (supported only on <b>mge</b> interfaces)</p> <p><b>5g</b>—5 Gbps (supported only on <b>mge</b> interfaces)</p> <p><b>10g</b>—10 Gbps (supported on <b>mge</b> interfaces and 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches)</p> | <p><i>speed</i>—Specify the interface speed.</p> <p><b>NOTE:</b> On 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module, no explicit configuration is required as it automatically identifies the transceivers and creates the interface accordingly.</p> |

**Required Privilege Level**

interface—To view this statement in the configuration.

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

**Release History Table**

| Release                | Description                                                                                                     |
|------------------------|-----------------------------------------------------------------------------------------------------------------|
| <a href="#">18.2R1</a> | Starting in Junos OS Release 18.1R2, the multi-rate speed is supported on EX2300-48MP and EX2300-24MP switches. |



## RELATED DOCUMENTATION

---

*Configuring the Interface Speed*

---

[Configuring the Interface Speed on Ethernet Interfaces | 7](#)

---

[Configuring Gigabit Ethernet Autonegotiation | 306](#)

---

*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](#)

---

[Configure Rate Selectability on ACX5448-D and ACX5448-M Routers | 417](#)

## speed (MX Series DPC)

### Syntax

```
speed (auto | 1Gbps | 100Mbps | 10Mbps);
```

### Hierarchy Level

```
[edit interfaces ge-fpc/pic/port]
```

### Release Information

Statement introduced in Junos OS Release 9.5.

### 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.

## RELATED DOCUMENTATION

[Configuring Gigabit Ethernet Autonegotiation](#) | 306

[no-auto-mdix](#) | 1285

## static-interface

### 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]
```

### Release Information

Statement introduced in Junos OS Release 10.2.

### 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.

## RELATED DOCUMENTATION

*Configuring PPPoE Service Name Tables*

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

## switch-options

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

## switch-port

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

## symbol-period

### 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]
```

### Release Information

Statement introduced in Junos OS Release 8.4.

### 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.

### RELATED DOCUMENTATION

[Configuring Threshold Values for Local Fault Events on an Interface](#) | 662

[Configuring Threshold Values for Fault Events in an Action Profile](#) | 677

## sync-reset

### 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]
```

### Release Information

Statement introduced in Junos OS Release 16.1R1.

### 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. By default, LACP handles the minimum-link failure. You can disable the minimum-link failure at LACP level, by using the **disable** option with the **sync-reset** command.

### Default

By default, LACP handles the minimum-link failure.

### Options

**disable**—To disable handling of the minimum-link failure at LACP level

**enable**—To enable handling of the minimum-link failure at LACP level

### Required Privilege Level

interface

## RELATED DOCUMENTATION

[minimum-links](#) | [1273](#)

[show lacp interfaces](#) | [1862](#)

## syslog (OAM Action)

### Syntax

```
syslog;
```

### Hierarchy Level

```
[edit protocols oam ethernet link-fault-management action-profile action]
```

### Release Information

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

Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

### Description

Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event.

Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.

### Required Privilege Level

interface—To view this statement in the configuration.

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

routing—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Specifying the Actions to Be Taken for Link-Fault Management Events](#) | 674

*Configuring Ethernet OAM Link Fault Management*



## system-id

### Syntax

```
system-id system-id;
```

### Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp]
```

### Release Information

Statement introduced in Junos OS Release 12.2R1

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### 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.

## RELATED DOCUMENTATION

[Configuring LACP for Aggregated Ethernet Interfaces](#) | 67

## system-priority

### Syntax

```
system-priority priority;
```

### Hierarchy Level

```
[edit interfaces aeX aggregated-ether-options lacp]
```

### Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Release 11.4 for EX Series switches.

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### 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.

## targeted-options (Grouping Subscribers by Bandwidth Usage)

### 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]
```

### 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.

### 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

**backup**—(Optional) Specify a backup member link per subscriber when you configure manual targeting.

**group**—(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 **default** group and distributed evenly across member links. Grouping of subscribers is supported only for static subscribers.

**Default:** default

**primary**—Specify a primary member link per subscriber when you configure manual targeting. You must always configure a primary link when you configure manual targeting.

**weight (\$junos-interface-target-weight | *weight-value*)**—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.

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.

**Range:** 1 through 1000

**Required Privilege Level**

interface—To view this statement in the configuration.

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

**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)

### 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 ]
```

### Release Information

Statement introduced in Junos OS Release 16.1.

### 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](#) 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.

RELATED DOCUMENTATION

|                                                                                             |
|---------------------------------------------------------------------------------------------|
| <a href="#">Ethernet Interfaces Overview   3</a>                                            |
| <a href="#">targeted-options (Grouping Subscribers by Bandwidth Usage)   1377</a>           |
| <i>Targeted Traffic Distribution on Aggregated Ethernet Interfaces in a Virtual Chassis</i> |

## targeted-distribution

### Syntax

```
targeted-distribution primary-list primary-list [backup-list backup-list;
```

### Hierarchy Level

```
[edit logical-systems name interfaces name unit ]
```

### Release Information

Statement introduced in Junos OS Release 16.1R1.

### 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.

## RELATED DOCUMENTATION

[Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 100](#)

[distribution-list | 1140](#)

[targeted-options | 1382](#)

## targeted-options

### 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]
```

### Release Information

Statement introduced in Junos OS Release 16.1.

### 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.

## RELATED DOCUMENTATION

[Targeted Distribution of Static Logical interfaces Across Aggregated Ethernet Member Links | 100](#)  
[distribution-list | 1140](#)  
[targeted-distribution | 1381](#)



## terminate (PPPoE Service Name Tables)

### 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]
```

### 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.

### 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.

## RELATED DOCUMENTATION

| *Configuring PPPoE Service Name Tables*

## thresholds

### 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]
```

### Release Information

Statement introduced in Junos OS Release 15.1.

### 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

**clear *clear-value***—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.

**interval *interval-value***—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.

**set *set-value***—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.

**warning clear *warning-clear-value***—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  $1\text{E}-3$  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.

**warning set warning-set-value**—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  $1\text{E}-n$  format, where 1 is the mantissa (remains constant) and  $n$  is the exponent. For example, a threshold value of  $1\text{E}-3$  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.

#### RELATED DOCUMENTATION

---

[Link Degrade Monitoring Overview | 843](#)

---

[link-degrade-monitor | 1227](#)

---

[recovery | 1335](#)

---

[request interface link-degrade-recover | 1476](#)

## traceoptions (LLDP)

### 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]
```

### Release Information

Statement introduced in Junos OS Release 9.6.

### Description

Set LLDP protocol-level tracing options.

### 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. 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**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the LLDP-specific tracing options:

- **all**—Trace all operations.
- **config**—Log configuration events.

- **interface**—Trace interface update events.
- **protocol**—Trace protocol information.
- **rtsock**—Trace socket events.
- **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.

RELATED DOCUMENTATION

| *Tracing LLDP Operations*

## traceoptions (Individual Interfaces)

### List of Syntax

[Syntax \(Individual interfaces with PTX Series, EX Series, ACX Series\) on page 1389](#)

[Syntax \(Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series\) on page 1389](#)

[Syntax \(OAMLFM with EX Series, QFX Series, NFX Series\) on page 1389](#)

[Syntax \(Interface process with ACX Series, SRX Series, MX Series, M Series, T Series\) on page 1389](#)

### 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]

### 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 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in JUNOS Release 10.2 for EX Series switches.

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

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

### Description

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 122 on page 1392](#) lists options for traceoption command for the following platforms:

Table 122: 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 filename</b> | <p>—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.</p>                                       |                                                                                    | <p>—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>.</p>                                                                                                                                                                                                                                                                                                         | <p>—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>.</p>                                                                                                                                                                                   |
| <b>files number</b>  | <p>—(Optional) 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>—(Optional) 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 <b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> 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) 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: 2 through</i></p> |

Table 122: 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                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 1000<br><i>Default: 3 files</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>flag</b>  | <p>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> 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> <li>• <b>q921</b>—Trace ISDN Q.921 frames</li> <li>• <b>q931</b>—Trace ISDN Q.931 frames</li> </ul> | <p>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> 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> <li>• <b>q921</b>—Trace ISDN Q.921 frames</li> <li>• <b>q931</b>—Trace ISDN Q.931 frames</li> </ul> | <p>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> 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 <b>flag</b> 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 configuration messages to kernel</li> </ul> |
| <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 122: 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 size | <p>—(Optional) 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) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). 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. If you specify a maximum number of files, you also must specify a maximum file size with the <b>files</b> option.</p> <p><b>Syntax:</b> <b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify GB</p> <p><b>Range:</b> 10 KB through 1 GB</p> <p><b>Default:</b> 128 KB</p> <p><b>Default:</b> If you do not include this option, tracing output is appended to an existing trace file.</p> |                                                                              |

Table 122: Options for traceoptions (continued)

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------|--------------------------------------------------------------|------------------------------------------------------------------------------------|----------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|        |                                                              |                                                                                    |                                                                      | <p>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <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 of trace files with the <b>files</b> option.</p> <p>Syntax: <b>xk</b> to specify kilobytes, <b>xm</b> to specify megabytes, or <b>xg</b> to specify gigabytes</p> <p>Range: 10 KB through the maximum file size</p> |

Table 122: 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                                                                                                     |
|--------------------------|--------------------------------------------------------------|------------------------------------------------------------------------------------|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                          |                                                              |                                                                                    |                                                                      | supported on your router<br><br><i>Default: 1 MB</i>                                                                                                                             |
| <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.                                                                                                                                 |
| <b>disable</b>           |                                                              |                                                                                    |                                                                      | —(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b> . |
| <b>no-remote-trace</b>   |                                                              |                                                                                    | —(Optional) Disable the remote trace.                                | -                                                                                                                                                                                |
| <b>match regex</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.

**RELATED DOCUMENTATION**

|                                                                       |
|-----------------------------------------------------------------------|
| <i>Tracing Operations of an Individual Router Interface</i>           |
| <i>Tracing Operations of an Individual Router or Switch Interface</i> |
| <i>Example: Configuring Ethernet OAM Link Fault Management</i>        |
| <i>Configuring Ethernet OAM Link Fault Management</i>                 |
| <i>Tracing Operations of the Interface Process</i>                    |

## traceoptions (LACP)

### Syntax

```
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

### Hierarchy Level

```
[edit protocols lacp]
```

### Release Information

Statement introduced in Junos OS Release 7.6.

Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

### 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.

#### RELATED DOCUMENTATION

| [Configuring LACP for Aggregated Ethernet Interfaces](#) | 67

## traceoptions (PPPoE)

### 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]
```

### Release Information

Statement introduced in Junos OS Release 9.6.

Option **filter** introduced in Junos OS Release 12.3

### 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.

#### **RELATED DOCUMENTATION**

*Configuring PPPoE Service Name Tables*

[Tracing PPPoE Operations](#) | 53

## tx-duration

### Syntax

```
expected-defect {  
  tx-duration ;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

### Release Information

Statement introduced in Junos OS Release 19.1.

### 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*.

**Maximum 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.

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | 1132

[show oam ethernet connectivity-fault-management mep-database](#) | 1901

## tx-enable

### Syntax

```
expected-defect {  
  tx-enable ;
```

### Hierarchy Level

```
[edit protocols oam ethernet connectivity-fault-management expected-defect]
```

### Release Information

Statement introduced in Junos OS Release 19.1.

### 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.

## RELATED DOCUMENTATION

[connectivity-fault-management](#) | 1132

[show oam ethernet connectivity-fault-management mep-database](#) | 1901

## underlying-interface

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

Support for aggregated Ethernet added in Junos OS Release 9.4.

### 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.

## RELATED DOCUMENTATION

*Configuring an IP Demultiplexing Interface*

---

*Configuring a VLAN Demultiplexing Interface*

---

[Configuring the PPPoE Underlying Interface | 48](#)

---

*Junos OS Interfaces and Routing Configuration Guide*



## unit

### 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;
            user-prefix user-prefix-string;
          }
        }
      }
      dynamic-profile profile-name {
        network ip-address {
          range name {
            low lower-limit;
            high upper-limit;
          }
        }
      }
    }
  }
}

inet6 {
  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;
          user-prefix user-prefix-string;
        }
      }
    }
    dynamic-profile profile-name {
      network ip-address {

```



```

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;
    master-only;
    multipoint-destination address {
        dlci dlcid-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;
    primary;
    (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;
        track {
            interface interface-name {
                bandwidth-threshold bits-per-second priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/prefix-length routing-instance instance-name priority-cost cost;
        }
        virtual-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {

```

```

        active-interface interface-name;
        active-group group-number;
    }
}
}
}
}

```

## Hierarchy Level

```

[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name],
[edit interfaces interface-set interface-set-name interface interface-name]

```

## 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.

## Description

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

## Options

***logical-unit-number***—Number of the logical unit.

**Range:** 0 through 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.

## RELATED DOCUMENTATION

*Configuring Logical Interface Properties*

*Junos OS Services Interfaces Library for Routing Devices*

## unnumbered-address (Dynamic Profiles)

### 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]
```

### 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.

### 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.

## RELATED DOCUMENTATION

| [Dynamic Profiles Overview](#)

## unnumbered-address (PPP)

### 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]
```

### Release Information

Statement introduced before Junos OS Release 7.4.

### 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.

## RELATED DOCUMENTATION

| [Configuring IPCP Options for Interfaces with PPP Encapsulation](#)

## Unused

### Syntax

```
unused;
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Evolved Release 19.3R1 for PTX10003-80C and PTX10003-160C routers.

### Description

, you cannot use port 1,3, 6 and/or 7 of the same logical PIC.

(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.

Refer to [“Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers” on page 381](#) for more information.

### Options

None

### Required Privilege Level

interface—To view this statement in the configuration.

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

## RELATED DOCUMENTATION

[Configuring 400-Gigabit Ethernet Interfaces on PTX10003 Routers | 381](#)

[speed \(Ethernet\) | 1359](#)

[number-of-sub-ports | 1291](#)

## virtual-control-channel

### 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)]
```

### Release Information

Statement introduced in Junos OS Release 14.2.

### 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.

## RELATED DOCUMENTATION

[Ethernet Ring Protection Switching Overview](#)

[Configuring Ethernet Ring Protection Switching on Switches \(CLI Procedure\)](#)



## virtual-switch

### 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]
```

### Release Information

Statement introduced in Junos OS Release 9.6.

### 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.

### RELATED DOCUMENTATION

| [Configuring Maintenance Intermediate Points \(MIPs\)](#) | 556

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

### Syntax

```
vlan-rule (high-low | odd-even);
```

### Hierarchy Level

```
[edit chassis fpc slot pic slot forwarding-mode vlan-steering]
```

### Release Information

Statement introduced in Junos OS Release 10.4.

### 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.

## RELATED DOCUMENTATION

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP](#) | 256

[forwarding-mode \(100-Gigabit Ethernet\)](#) | 1189

[vlan-steering \(100-Gigabit Ethernet Type 4 PIC with CFP\)](#) | 1425

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

### Syntax

```
vlan-steering {  
    vlan-rule (high-low | odd-even);  
}
```

### Hierarchy Level

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

### Release Information

Statement introduced in Junos OS Release 9.4.

### 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.

## RELATED DOCUMENTATION

[Configuring VLAN Steering Mode for 100-Gigabit Ethernet Type 4 PIC with CFP | 256](#)

[forwarding-mode \(100-Gigabit Ethernet\) | 1189](#)

[sa-multicast \(100-Gigabit Ethernet\) | 1344](#)

[vlan-rule \(100-Gigabit Ethernet Type 4 PIC with CFP\) | 1424](#)

# Operational Commands

## IN THIS CHAPTER

- clear interfaces interface-set statistics | 1430
- clear interfaces interval | 1431
- clear interfaces aeX forwarding-options load-balance state | 1433
- clear interfaces aggregate forwarding-options load-balance state | 1434
- clear interfaces transport pm | 1435
- clear oam ethernet connectivity-fault-management continuity-measurement | 1437
- clear oam ethernet connectivity-fault-management linktrace path-database | 1439
- clear oam ethernet connectivity-fault-management loss-statistics | 1440
- clear oam ethernet connectivity-fault-management policer | 1442
- clear oam ethernet connectivity-fault-management statistics | 1444
- clear oam ethernet connectivity-fault-management synthetic-loss-measurement | 1446
- clear oam ethernet link-fault-management state | 1448
- clear oam ethernet link-fault-management statistics | 1449
- clear protection-group ethernet-ring statistics | 1450
- monitor ethernet delay-measurement | 1451
- monitor ethernet loss-measurement | 1457
- monitor ethernet synthetic-loss-measurement | 1463
- monitor ethernet synthetic-loss-measurement | 1468
- prbs-test-start | 1473
- prbs-test-stop | 1475
- request interface link-degrade-recover | 1476
- request interface mc-ae switchover (Multichassis Link Aggregation) | 1480
- request interface (revert | switchover) (Aggregated Ethernet Link Protection) | 1482
- request lacp link-switchover | 1484
- show chassis hardware | 1485
- show chassis pic | 1504
- show ethernet-switching redundancy-groups | 1540
- show interfaces (Adaptive Services) | 1545

- [show interfaces \(Aggregated Ethernet\) | 1553](#)
- [show interfaces demux0 \(Demux Interfaces\) | 1568](#)
- [show interfaces diagnostics optics \(Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Virtual Chassis Port\) | 1580](#)
- [show interfaces \(far-end-interval\) | 1631](#)
- [show interfaces \(Fast Ethernet\) | 1633](#)
- [show interfaces | 1656](#)
- [show interfaces \(M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet\) | 1763](#)
- [show interfaces \(PPPoE\) | 1786](#)
- [show interfaces interface-set \(Ethernet Interface Set\) | 1799](#)
- [show interfaces interface-set queue | 1805](#)
- [show interfaces interval | 1816](#)
- [show interfaces irb | 1820](#)
- [show interfaces mac-database | 1828](#)
- [show interfaces mc-ae | 1835](#)
- [show interfaces prbs-stats | 1842](#)
- [show interfaces transport pm | 1844](#)
- [show l2-learning instance | 1854](#)
- [show l2-learning redundancy-groups | 1856](#)
- [show lacp interfaces | 1862](#)
- [show oam ethernet connectivity-fault-management delay-statistics | 1869](#)
- [show oam ethernet connectivity-fault-management forwarding-state | 1874](#)
- [show oam ethernet connectivity-fault-management interfaces | 1879](#)
- [show oam ethernet connectivity-fault-management linktrace path-database | 1893](#)
- [show oam ethernet connectivity-fault-management loss-statistics | 1896](#)
- [show oam ethernet connectivity-fault-management mep-database | 1901](#)
- [show oam ethernet connectivity-fault-management mep-statistics | 1916](#)
- [show oam ethernet connectivity-fault-management path-database | 1932](#)
- [show oam ethernet connectivity-fault-management policer | 1934](#)
- [show oam ethernet connectivity-fault-management sla-iterator-statistics | 1938](#)
- [show oam ethernet connectivity-fault-management synthetic-loss-statistics | 1951](#)
- [show oam ethernet evc | 1954](#)
- [show oam ethernet fnp interface | 1956](#)
- [show oam ethernet fnp messages | 1958](#)

- [show oam ethernet fnp status | 1960](#)
- [show oam ethernet link-fault-management | 1962](#)
- [show oam ethernet lmi | 1972](#)
- [show oam ethernet lmi statistics | 1974](#)
- [show pppoe interfaces | 1976](#)
- [show pppoe service-name-tables | 1981](#)
- [show pppoe sessions | 1985](#)
- [show pppoe statistics | 1987](#)
- [show pppoe underlying-interfaces | 1990](#)
- [show pppoe version | 1999](#)
- [show protection-group ethernet-ring aps | 2001](#)
- [show protection-group ethernet-ring configuration | 2006](#)
- [show protection-group ethernet-ring data-channel | 2015](#)
- [show protection-group ethernet-ring flush-info | 2018](#)
- [show protection-group ethernet-ring interface | 2020](#)
- [show protection-group ethernet-ring node-state | 2025](#)
- [show protection-group ethernet-ring statistics | 2031](#)
- [show protection-group ethernet-ring vlan | 2038](#)
- [traceroute ethernet | 2044](#)

## clear interfaces interface-set statistics

### Syntax

```
clear interfaces interface-set statistics interface-set-name
```

### Release Information

Command introduced in Junos OS Release 8.5.

### 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

### List of Sample Output

[clear interfaces interface-set statistics on page 1430](#)

### 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
```

## clear interfaces interval

### Syntax

```
clear interfaces interval interface-name
```

### Release Information

Command introduced before Junos OS Release 7.4.

### 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

## RELATED DOCUMENTATION

[show interfaces interval](#) | 1816

### List of Sample Output

[clear interfaces interval on page 1431](#)

### Output Fields

See [show interfaces interval](#) 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

```

## clear interfaces aeX forwarding-options load-balance state

### Syntax

```
clear interfaces aeX unit logical-unit-number aggregate forwarding-options load-balance state
```

### Release Information

Command introduced in Junos OS Release 13.2R1.

### 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

**aeX**—Name of a particular aggregated Ethernet interface.

***logical-unit-number***—Number of the logical unit of the interface.

**forwarding-options load-balance state**—Cause the load-balancing state to be cleared for the specific interface.

### Required Privilege Level

clear

### RELATED DOCUMENTATION

| [show interfaces interval](#) | 1816

### List of Sample Output

[clear interfaces aeX aggregate forwarding-options on page 1433](#)

## Sample Output

```
clear interfaces aeX aggregate forwarding-options
```

```
user@host> clear interfaces ae1 aggregate forwarding-options load-balance state
```

## clear interfaces aggregate forwarding-options load-balance state

### Syntax

```
clear interfaces aggregate forwarding-options load-balance state
```

### Release Information

Command introduced in Junos OS Release 13.2R1.

### 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

### RELATED DOCUMENTATION

| [show interfaces interval](#) | [1816](#)

### List of Sample Output

[clear interfaces aggregate forwarding-options on page 1434](#)

## Sample Output

```
clear interfaces aggregate forwarding-options
```

```
user@host> clear interfaces aggregate forwarding-options load-balance state
```

## clear interfaces transport pm

### Syntax

```
clear interfaces transport pm (all | optics | otn) (all | current | current-day) (all | interface-name)
```

### Release Information

Command introduced in Junos OS Release 14.2 on the PTX Series.

Command introduced in Junos OS release 16.1 on the MX Series.

Command introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

### 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

### RELATED DOCUMENTATION

---

[show interfaces transport pm | 1844](#)

---

[100-Gigabit Ethernet OTN Options Configuration Overview | 424](#)

---

[tca | 972](#)

### List of Sample Output

[clear interfaces transport pm on page 1436](#)

### 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
```

## clear oam ethernet connectivity-fault-management continuity-measurement

### Syntax

```
clear oam ethernet connectivity-fault-management continuity-measurement
maintenance-domain md-name
maintenance-association ma-name
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

### Release Information

Command introduced in Junos OS Release 11.1.

### Description

For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear the existing continuity measurement and restart counting the operational uptime (that is, the total time during which CCM adjacency is active for a particular remote MEP.).

### Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**local-mep *local-mep-id***—(Optional) Display connectivity fault management information for the specified local MEP only.

**remote-mep *remote-mep-id***—(Optional) Display connectivity fault management information for the specified remote MEP only.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[Managing Continuity Measurement Statistics | 740](#)

*Ethernet Interfaces User Guide for Routing Devices*

### List of Sample Output

[clear oam ethernet connectivity-fault-management continuity-measurement on page 1438](#)

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

```
clear oam ethernet connectivity-fault-  
management continuity-measurement
```

```
user@host> clear oam ethernet connectivity-fault-management continuity-measurement  
maintenance-domain md5 maintenance-association ma5 local-mep 100 remote-mep 102
```

```
Continuity measurement restarted.
```

## clear oam ethernet connectivity-fault-management linktrace path-database

### Syntax

```
clear oam ethernet connectivity-fault-management linktrace path-database mac-address maintenance-association
    ma-name maintenance-domain md-name
```

### Release Information

Command introduced in Junos OS Release 9.0.

### Description

Clear all the linktrace entries and the relevant path information from the database for a particular remote host on M320, MX Series, T320, and T640 routers.

### Options

**mac-address**—Clear connectivity fault management path database information for the specified MAC address of the remote host.

**maintenance-association *ma-name***—Clear connectivity fault management path database information for the specified maintenance association.

**maintenance-domain *md-name***—Clear connectivity fault management path database information for the specified maintenance domain.

### Required Privilege Level

view

## Sample Output

```
clear oam ethernet connectivity-fault-
management linktrace path-database
```

```
user@host> clear oam ethernet connectivity-fault-management linktrace path-database
maintenance-domain md1 maintenance-association ma3 00058573e483
```

This command produces no output.



## clear oam ethernet connectivity-fault-management loss-statistics

### Syntax

```
clear oam ethernet connectivity-fault-management loss-statistics  
<interface ethernet-interface-name>  
<level md-level>
```

### Release Information

Command introduced in Junos OS Release 11.1.

### Description

For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear all loss statistics maintained by CFM for a given maintenance domain and maintenance association.

In addition, for Ethernet interfaces on MX Series routers, clear any ITU-T Y.1731 Ethernet frame loss measurement (ETH-LM) statistics.

By default, the command clears ETH-LM statistics for CFM maintenance association end points (MEPs) attached to any interface on the router.

### Options

**interface *ethernet-interface-name***—(Optional) Clear ETH-LM statistics and ETH-LM frame counts only for MEPs attached to the specified Ethernet physical interface.

**level *md-level***—(Optional) Clear ETH-LM statistics and ETH-LM frame counts only for MEPs within CFM maintenance domains (MDs) of the specified level.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[Clearing ETH-LM Statistics | 750](#)

[Displaying ETH-LM Statistics | 749](#)

[Managing ETH-LM Statistics | 748](#)

### List of Sample Output

[clear oam ethernet connectivity-fault-management loss-statistics on page 1441](#)

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

```
clear oam ethernet connectivity-fault-  
management loss-statistics
```

```
user@host> clear oam ethernet connectivity-fault-management loss-statistics
```

```
Cleared loss measurements statistics of all CFM sessions
```

## clear oam ethernet connectivity-fault-management policer

### Syntax

```
clear oam ethernet connectivity-fault-management policer maintenance-domain md-name maintenance-association ma-name
```

### Release Information

Command introduced in Junos OS Release 10.0.

### Description

On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, M120, MX Series, T320, and T640 routers, clear connectivity-fault-management policer statistics.

### Options

The following options are supported:

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain. If this option is not specified, policer statistics are cleared for all maintenance associations for all maintenance domains.

**maintenance-association *ma-name*** —Name of an existing CFM maintenance association. If this option is not specified, policer statistics are cleared for all maintenance associations for given maintenance domain. This option cannot be specified without specifying maintenance-domain name.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[show oam ethernet connectivity-fault-management policer](#) | 1934

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## clear oam ethernet connectivity-fault-management policer

```
user@host> clear oam ethernet connectivity-fault-management policer
```

```
Policer statistics cleared
```

**clear oam ethernet connectivity-  
fault-management policer maintenance-domain md-name  
maintenance-association ma-name**

```
user@host> clear oam ethernet connectivity-fault-management policer  
maintenance-domain md5 maintenance-association ma5-1
```

```
Policer statistics cleared
```

## clear oam ethernet connectivity-fault-management statistics

### Syntax

```
clear oam ethernet connectivity-fault-management statistics
<interface ethernet-interface-name>
<level md-level>
```

### Release Information

Command introduced in Junos OS Release 8.4.

Support for ETH-DM statistics and frame counts added in Junos OS Release 9.5.

### Description

For all routers that support IEEE 802.1ag OAM connectivity-fault management (CFM), clear all statistics maintained by CFM.

In addition, for Ethernet interfaces on Dense Port Concentrators (DPCs) in MX Series routers only, also clear any ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) statistics and ETH-DM frame counts.

By default, the command clears CFM statistics and ETH-DM statistics and frame counts for CFM maintenance association end points (MEPs) attached to any interface on the router.

### Options

**ethernet-interface-name**—(Optional) Clear CFM statistics, ETH-DM statistics, and ETH-DM frame counts only for MEPs attached to the specified Ethernet physical interface.

**level**—(Optional) Clear CFM statistics, ETH-DM statistics, and ETH-DM frame counts only for MEPs within CFM maintenance domains (MDs) of the specified level.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[show oam ethernet connectivity-fault-management delay-statistics](#) | 1869

[show oam ethernet connectivity-fault-management interfaces](#) | 1879

[show oam ethernet connectivity-fault-management mep-database](#) | 1901

[show oam ethernet connectivity-fault-management mep-statistics](#) | 1916

### List of Sample Output

[clear oam ethernet connectivity-fault-management statistics on page 1445](#)

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**clear oam ethernet connectivity-fault-management statistics**

**user@host> clear oam ethernet connectivity-fault-management statistics**

Cleared statistics of all CFM sessions

## clear oam ethernet connectivity-fault-management synthetic-loss-measurement

### Syntax

```
clear oam ethernet connectivity-fault-management synthetic-loss-measurement
maintenance-domain md-name
maintenance-association ma-name
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

### Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.

### Description

(MX Series routers)—For all routers that support IEEE 802.1ag OAM connectivity fault management (CFM), clear the existing on-demand Ethernet synthetic loss measurement (ETH-SLM) statistics and restart counting the ETH-SLM frame counts and statistics.

### Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**local-mep *local-mep-id***—(Optional) Clear connectivity fault management information for the specified local MEP only.

**remote-mep *remote-mep-id***—(Optional) Clear connectivity fault management information for the specified remote MEP only.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[monitor ethernet synthetic-loss-measurement | 1463](#)

[show oam ethernet connectivity-fault-management synthetic-loss-statistics | 1951](#)

### List of Sample Output

[clear oam ethernet connectivity-fault-management synthetic-loss-measurement on page 1447](#)

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

```
clear oam ethernet connectivity-fault-  
management synthetic-loss-  
measurement
```

```
user@host> clear oam ethernet connectivity-fault-management synthetic-loss-measurement  
maintenance-domain md5 maintenance-association ma5 local-mep 100 remote-mep 102
```

```
Synthetic loss measurement restarted.
```



## clear oam ethernet link-fault-management state

### Syntax

```
clear oam ethernet link-fault-management state <interface-name>
```

### Release Information

Command introduced in Junos OS Release 8.4.

### Description

On all M Series, MX Series, ACX series, PTX Series, T320, and T640 routers, clear link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on Ethernet interfaces.

### Options

**none**—Clear OAM link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on all Ethernet interfaces.

**interface-name**—(Optional) Clear OAM link fault management state information, restart the link discovery process, and reset OAM loopback state (if set previously) on the specified Ethernet interface only.

### Required Privilege Level

view

### List of Sample Output

[clear oam ethernet link-fault-management state on page 1448](#)

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

```
clear oam ethernet link-fault-management state
```

```
user@host> clear oam ethernet link-fault-management state ge-0/3/3
```

```
Cleared link-fault-management state for interface ge-0/3/3
```

## clear oam ethernet link-fault-management statistics

### Syntax

```
clear oam ethernet link-fault-management <interface-name>
```

### Release Information

Command introduced in Junos OS Release 8.2.

### Description

On M320, M120, MX Series, PTX Series, T320, and T640 routers, clear Operation, Administration, and Management (OAM) link fault management statistics or state information from Ethernet interfaces.

### Options

**none**—Clear OAM link fault management statistics from all Ethernet interfaces.

**interface-name**—(Optional) Clear OAM link fault management statistics from the specified Ethernet interface only.

### Required Privilege Level

view

### List of Sample Output

[clear oam ethernet link-fault-management statistics on page 1449](#)

### Output Fields

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

```
clear oam ethernet link-fault-management statistics
```

```
user@host> clear oam ethernet link-fault-management statistics
```

```
Cleared link-fault-management statistics for all interfaces
```

## clear protection-group ethernet-ring statistics

### Syntax

```
clear protection-group ethernet-ring statistics  
<group-name group-name>
```

### Release Information

Command introduced in Junos OS Release 9.4.

### 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

### List of Sample Output

[clear protection-group ethernet-ring statistics on page 1450](#)

[clear protection-group ethernet-ring statistics on page 1450](#)

### 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
```

## monitor ethernet delay-measurement

### Syntax

```
monitor ethernet delay-measurement
maintenance-domain md-name
maintenance-association ma-name
(one-way | two-way)
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<local-mep mep-id>
<wait interval-seconds>
<priority 802.1p value>
<size>
<no-session-id-tlv>
<xml>
```

### Release Information

Command introduced in Junos OS Release 9.5.

**local-mep** option introduced in Junos OS Release 15.1.

### Description

Start an ITU-T Y.1731 Ethernet frame delay measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. Frame delay measurement statistics are stored at one of the MEPs for later retrieval.

**NOTE:** If you attempt to monitor delays to a nonexistent MAC address, you must type Ctrl +c to explicitly quit the **monitor ethernet delay-measurement** command and return to the CLI command prompt.

To start an Ethernet frame delay measurement session, the router initiates an exchange of frames carrying one-way or two-way frame delay measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet frame delay measurement PDU frames exchanged to measure frame delay times—are displayed as the runtime output of the **monitor ethernet delay-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet frame delay measurement statistics, described below, are measured and stored at only one of the MEPs:

**Frame delay**—The difference, in microseconds, between the time a frame is sent and when it is received.

**Frame delay variation**—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

## Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**one-way**—Measurement type is one-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a one-way delay measurement request (1DM) frame and the time at which the receiver MEP receives the frame.

**two-way**—Measurement type is two-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a two-way delay measurement message (DMM) frame and the time at which the initiator MEP receives an associated two-way delay measurement reply (DMR) frame from the responder MEP, subtracting the time elapsed at the responder MEP.

**mep *remote-mep-id***—Numeric identifier of the peer MEP with which to perform Ethernet frame delay measurement. The discovered MAC address of the peer MEP is used. The range of values is 1 through 8191.

**remote-mac-address**—Unicast MAC address of the peer MEP with which to perform Ethernet frame delay measurement. 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**. Multicast MAC addresses are not supported.

**count *frame-count***—(Optional) Number of frames to send to the specified peer MEP. The range of values is 1 through 65,535 frames. The default value is 10 frames.

**local-mep *mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance association end point.

**wait *interval-seconds***—(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.

**priority *802.1p value***—(Optional) Priority of the delay measurement request frame supported by both one-way delay measurement and two-way delay measurement. The range of values is from 0 through 7. The default value is zero.

**size**—(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

**no-session-id-tlv**—(Optional) Prevent insertion of the session ID TLV in the request frame.

**xml**—(Optional) Allow the output of the command to be displayed in XML format supported by both one-way delay measurement and two-way delay measurement. Note that the only way to get output in XML format is to use the **xml** argument. The **display xml** command does not work.

### Additional Information

To display the frame counts collected at an MEP as the result of this command, see the following command descriptions in the [CLI Explorer](#):

- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at an MEP as the result of this command, see the following command descriptions in the [CLI Explorer](#).

- **show oam ethernet connectivity-fault-management delay-statistics**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management statistics** command, described in the [CLI Explorer](#).

For a complete description of Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.

### Required Privilege Level

trace and maintenance

### List of Sample Output

[monitor ethernet delay-measurement one-way on page 1455](#)

[monitor ethernet delay-measurement two-way on page 1455](#)

[monitor ethernet delay-measurement two-way \(Invalid DMR Frames Received\) on page 1456](#)

### Output Fields

The **monitor ethernet delay-measurement** command displays different output at the CLI, depending on whether you start a one-way or two-way frame delay measurement:

- [Table 123 on page 1454](#) lists the run-time output fields for the **monitor ethernet delay-measurement one-way** command.
- [Table 124 on page 1454](#) lists the run-time output fields for the **monitor ethernet delay-measurement two-way** command.

Output fields are listed in the approximate order in which they appear.

Table 123: monitor ethernet delay-measurement one-way Output Fields

| Output Field Name                | Output Field Description                                                                       |
|----------------------------------|------------------------------------------------------------------------------------------------|
| <b>One-way ETH-DM request to</b> | Unicast MAC address of the remote peer MEP.                                                    |
| <b>Interface</b>                 | Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached. |
| <b>1DM Frames sent</b>           | PDU frames sent to the remote MEP in this ETH-DM session.                                      |
| <b>Packets transmitted</b>       | Total number of 1DM PDU frames sent to the remote MEP during this measurement session.         |
| <b>Average delay</b>             | Average two-way frame delay measured in this session.                                          |
| <b>Average delay variation</b>   | Average frame jitter measured in this session.                                                 |
| <b>Best case delay</b>           | Lowest two-way frame delay measured in this session.                                           |
| <b>Worst case delay</b>          | Highest two-way frame delay measured in this session.                                          |

**NOTE:** For one-way delay measurement, these CLI output fields display **NA** ("not applicable") at the initiator MEP because one-way frame delay measurements occur at the receiver MEP.

Table 124: monitor ethernet delay-measurement two-way Output Fields

| Output Field Name                                          | Output Field Description                                                                                        |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| <b>Two-way Ethernet frame delay measurement request to</b> | Unicast MAC address of the remote peer MEP.                                                                     |
| <b>Interface</b>                                           | Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.                  |
| <b>DMR received from</b>                                   | Unicast MAC address of the remote MEP that transmitted this DMR frame in response to a DMM frame.               |
| <b>Delay</b>                                               | Two-way delay, in microseconds, for the initiator-transmitted DMM frame.                                        |
| <b>Delay variation</b>                                     | Difference, in microseconds, between the current and previous delay values. This is also known as frame jitter. |
| <b>Packets transmitted</b>                                 | Total number of DMM PDU frames sent to the remote MEP in this measurement session.                              |

Table 124: monitor ethernet delay-measurement two-way Output Fields (*continued*)

| Output Field Name              | Output Field Description                                                                 |
|--------------------------------|------------------------------------------------------------------------------------------|
| <b>Valid packets received</b>  | Total number of DMR PDU frames received from the remote MEP in this measurement session. |
| <b>Average delay</b>           | Average two-way frame delay measured in this session.                                    |
| <b>Average delay variation</b> | Average frame jitter measured in this session.                                           |
| <b>Best case delay</b>         | Lowest two-way frame delay measured in this session.                                     |
| <b>Worst case delay</b>        | Highest two-way frame delay measured in this session.                                    |

## Sample Output

### monitor ethernet delay-measurement one-way

```
user@host> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```

```
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```

### monitor ethernet delay-measurement two-way

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```

```
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
```



```
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec
```

```
--- Delay measurement statistics ---
```

```
Packets transmitted: 10, Valid packets received: 10
```

```
Average delay: 103 usec, Average delay variation: 8 usec
```

```
Best case delay: 92 usec, Worst case delay: 122 usec
```

### monitor ethernet delay-measurement two-way (Invalid DMR Frames Received)

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```

```
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
```

```
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a with invalid timestamp(s).
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec
```

```
--- Delay measurement statistics ---
```

```
Packets transmitted: 10, Valid packets received: 9, Invalid packets received: 1
```

```
Average delay: 105 usec, Average delay variation: 9 usec
```

```
Best case delay: 92 usec, Worst case delay: 122 usec
```

## monitor ethernet loss-measurement

### Syntax

```
monitor ethernet loss-measurement
maintenance-domain md-name
maintenance-association ma-name
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<local-mep mep-id>
<wait interval-seconds>
<priority 802.1p value>
<no-session-id-tlv>
<xml>
```

### Release Information

Command introduced in Junos OS Release 11.1.

**local-mep** option introduced in Junos OS Release 15.1

### Description

Start an ITU-T Y.1731 Ethernet frame loss measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a count of transmitted and received data frames between the pair of MEPs. Frame loss measurement statistics are stored at one of the MEPs for later retrieval. For MX Series routers, supports point-to-point down MEPs for Ethernet interfaces (as per IEEE 802.1ag over VPWS).

**NOTE:** If you attempt to monitor loss to a nonexistent MAC address, you must type Ctrl + c to explicitly quit the **monitor ethernet loss-measurement** command and return to the CLI command prompt.

To start an Ethernet frame loss measurement session, the router first sends frames with ETH-LM information to a peer MEP and similarly receives frames with ETH-LM information from the peer MEP. Frame loss is calculated by collecting the counter values applicable for ingress and egress service frames where the counters maintain a count of transmitted and received data frames between a pair of MEPs. The loss measurement statistics are retrieved as the output of the **monitor ethernet loss-measurement** command and are also stored at the initiator. The frames counts are stored at both the initiator and the receiver MEPs for later retrieval.

### Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**mep remote-mep-id**—Numeric identifier of the peer MEP with which to perform Ethernet frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8192.

**remote-mac-address**—Unicast MAC address of the peer MEP with which to perform Ethernet frame loss measurement. 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**). Multicast MAC addresses are not supported.

**count frame-count**—(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65535 frames. The default value is 10 frames.

**local-mep mep-id**—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

**wait interval-seconds**—(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.

**priority 802.1p value**—(Optional) Priority of the delay measurement request frame. The range of values is from 0 through 7. The default value is 1 second.

**no-session-id-tlv**—(Optional) Disable the **session id TLV** argument set in the request frame.

**xml**—(Optional) Allow the output of the command to be displayed in XML format.

### Additional Information

To display the iterator output for an LM session, run the following command:

- **show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator <profile> maintenance-association <MA> maintenance-domain <MD> local-mep <MEP> remote-mep <RMEP>**

To display the frame counts collected at an MEP as the result of this command, see the following command descriptions in the [CLI Explorer](#):

- **show oam ethernet connectivity-fault-management loss-statistics**
- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at an MEP as the result of this command, see the following command descriptions in the [CLI Explorer](#):

- **show oam ethernet connectivity-fault-management delay-statistics**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management loss-statistics maintenance-domain *md-name* maintenance-association *ma-name*** command, as described in the [CLI Explorer](#).

For a complete description of Ethernet frame loss measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.

### Required Privilege Level

trace and maintenance

## RELATED DOCUMENTATION

[Ethernet Frame Loss Measurement Overview | 696](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

[CLI Explorer](#)

### List of Sample Output

[monitor ethernet loss-measurement on page 1461](#)

### Output Fields

[Table 125 on page 1459](#) lists the output fields for the **monitor ethernet loss-measurement** command and their descriptions. Output fields are listed in the approximate order in which they appear.

Table 125: monitor ethernet loss-measurement output fields

| Output Field Name                                 | Output Field Description                                                                                                                                                       |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ethernet loss delay measurement request to</b> | Unicast MAC address of the remote peer MEP.                                                                                                                                    |
| <b>Interface</b>                                  | Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.                                                                                 |
| <b>LMR received from</b>                          | Unicast MAC address of the remote MEP that transmitted this LMR frame in response to a loss measurement message (LMM) frame.                                                   |
| <b>Near-end frame loss</b>                        | Count of frame loss associated with ingress data frames.                                                                                                                       |
| <b>Far-end frame loss</b>                         | Count of frame loss associated with egress data frames.                                                                                                                        |
| <b>Near-end loss ratio</b>                        | Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the ingress interface. |
| <b>Far-end loss ratio</b>                         | Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the egress interface.  |

Table 125: monitor ethernet loss-measurement output fields (*continued*)

| Output Field Name                     | Output Field Description                                                                 |
|---------------------------------------|------------------------------------------------------------------------------------------|
| <b>LMM packets transmitted</b>        | Total number of LMM PDU frames sent to the remote MEP in this measurement session.       |
| <b>LMR packets received</b>           | Total number of LMR PDU frames received from the remote MEP in this measurement session. |
| <b>Average near-end frame loss</b>    | Average frame loss measured in this session associated with ingress data frames.         |
| <b>Average near-end loss ratio</b>    | Average frame loss ratio measured in this session associated with ingress data frames.   |
| <b>Average far-end frame loss</b>     | Average frame loss measured in this session associated with egress data frames.          |
| <b>Average far-end loss ratio</b>     | Average frame loss ratio measured in this session associated with egress data frames.    |
| <b>Near-end best case frame loss</b>  | Lowest frame loss measured in this session associated with ingress data frames.          |
| <b>Near-end best case loss ratio</b>  | Lowest frame loss ratio measured in this session associated with ingress data frames.    |
| <b>Near-end worst case frame loss</b> | Highest frame loss measured in this session associated with ingress data frames.         |
| <b>Near-end worst case loss ratio</b> | Highest frame loss ratio measured in this session associated with ingress data frames.   |
| <b>Far-end best case frame loss</b>   | Lowest frame loss measured in this session associated with egress data frames.           |
| <b>Far-end best case loss ratio</b>   | Lowest frame loss ratio measured in this session associated with egress data frames.     |
| <b>Far-end worst case frame loss</b>  | Highest frame loss measured in this session associated with egress data frames.          |
| <b>Far-end worst case loss ratio</b>  | Highest frame loss ratio measured in this session associated with egress data frames.    |

Note that in the preceding table, the term *number of service frames not delivered* is the difference between the number of service frames arriving at the ingress Ethernet flow point and the number of service frames delivered at the egress Ethernet flow point in a point-to-point Ethernet connection.

## Sample Output

### monitor ethernet loss-measurement

```
user@host> monitor ethernet loss-measurement mep 2 64:87:88:6a:da:94 maintenance-domain md
maintenance-association ma count 10
```

```
ETH-LM request to 64:87:88:6a:da:94, Interface ge-2/3/2.0
```

```
LMR received from 64:87:88:6a:da:94
```

```
Near-end frame loss(CIR)      :0          Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0          Far-end frame loss(EIR):260
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.43537%
```

```
LMR received from 64:87:88:6a:da:94
```

```
Near-end frame loss(CIR)      :0          Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51546%
Near-end frame loss(EIR)      :0          Far-end frame loss(EIR):257
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.31615%
```

```
LMR received from 64:87:88:6a:da:94
```

```
Near-end frame loss(CIR)      :0          Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0          Far-end frame loss(EIR):261
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.77551%
```

```
LMR received from 64:87:88:6a:da:94
```

```
Near-end frame loss(CIR)      :0          Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0          Far-end frame loss(EIR):260
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.43537%
```

```
LMR received from 64:87:88:6a:da:94
```

```
Near-end frame loss(CIR)      :0          Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51020%
Near-end frame loss(EIR)      :0          Far-end frame loss(EIR):259
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.09524%
```

```

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0          Far-end frame loss(CIR):0
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.00000%
Near-end frame loss(EIR)      :0          Far-end frame loss(EIR):519
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.71795%

```

```

LMR received from 64:87:88:6a:da:94
Near-end frame loss(CIR)      :0          Far-end frame loss(CIR):1
Near-end frame loss ratio(CIR):0.00000% Far-end frame loss ratio(CIR):0.51020%
Near-end frame loss(EIR)      :0          Far-end frame loss(EIR):259
Near-end frame loss ratio(EIR):0.00000% Far-end frame loss ratio(EIR):88.09524%

```

```

--- Loss measurement statistics ---

```

```

LMM packets transmitted: 10, Valid LMR packets received: 8

```

```

Average near-end loss(CIR)      : 0.00000
Average near-end loss ratio(CIR) : 0.00000%
Average far-end loss(CIR)       : 0.42857
Average far-end loss ratio(CIR)  : 0.21941%
Near-end best case loss(CIR)     : 0
Near-end best case loss ratio(CIR) : 0.00000%
Near-end worst case loss(CIR)    : 0
Near-end worst case loss ratio(CIR): 0.00000%
Far-end best case loss(CIR)      : 0
Far-end best case loss ratio(CIR) : 0.00000%
Far-end worst case loss(CIR)     : 1
Far-end worst case loss ratio(CIR) : 0.51546%
Average near-end loss(EIR)       : 0.00000
Average near-end loss ratio(EIR)  : 0.00000%
Average far-end loss(EIR)        : 296.42857
Average far-end loss ratio(EIR)   : 88.41011%
Near-end best case loss(EIR)     : 0
Near-end best case loss ratio(EIR) : 0.00000%
Near-end worst case loss(EIR)    : 0
Near-end worst case loss ratio(EIR): 0.00000%
Far-end best case loss(EIR)      : 257
Far-end best case loss ratio(EIR) : 88.09524%
Far-end worst case loss(EIR)     : 519
Far-end worst case loss ratio(EIR) : 88.77551%

```

## monitor ethernet synthetic-loss-measurement

### Syntax

```
monitor ethernet synthetic-loss-measurement
maintenance-domain md-name
maintenance-association ma-name
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<local-mep mep-id>
<wait interval-milliseconds>
<priority 802.1p value>
<size>
<xml>
```

### Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.

**local-mep** option introduced in Junos OS Release 15.1

### Description

(MX Series routers) Start an ITU-T Y.1731 Ethernet synthetic loss measurement (ETH-SLM) session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. ETH-SLM (also called synthetic frame loss measurement) statistics are stored at one of the MEPs for later retrieval.

**NOTE:** If you attempt to monitor delays to a nonexistent MAC address, you must press Ctrl +c to explicitly quit the **monitor ethernet synthetic-loss-measurement** command and return to the CLI command prompt.



To start an Ethernet synthetic frame loss measurement session, the router initiates an exchange of frames carrying synthetic frame loss measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet synthetic frame loss measurement PDU frames exchanged to measure frame delay times—are displayed as the run-time output of the **monitor ethernet synthetic-loss-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet synthetic frame loss measurement statistics, described below, are measured and stored at only one of the MEPs:

**Frame delay**—The difference, in microseconds, between the time a frame is sent and when it is received.

**Frame delay variation**—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

### Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**mep *remote-mep-id***—Numeric identifier of the peer MEP with which to perform Ethernet synthetic frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8191.

***remote-mac-address***—Unicast MAC address of the peer MEP with which to perform Ethernet synthetic frame loss measurement. 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**. Multicast MAC addresses are not supported.

**count *frame-count***—(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65,535 frames. The default value is 10 frames.

**local-mep *mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

**wait *interval-milliseconds***—(Optional) Number of milliseconds to wait between sending frames. You must specify this value in multiples of 100 milliseconds. The range of values is from 100 through 50000 milliseconds. The default value is 100 milliseconds.

**priority *802.1p value***—(Optional) Priority of the ETH-SLM request frame supported. The range of values is from 0 through 7. The default value is zero.

**size**—(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

**xml**—(Optional) Allow the output of the command to be displayed in XML format for ETH-SLM. Note that the only way to get output in XML format is to use the **xml** argument. The **display xml** command does not work.

### Additional Information

To display the frame counts collected at a MEP as the result of this command, use the following commands as described in the [CLI Explorer](#):

- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at a MEP as the result of this command, use the following commands as described in the [CLI Explorer](#).

- **show oam ethernet connectivity-fault-management synthetic-loss-measurement**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management statistics** command, described in the [CLI Explorer](#).

For a complete description of Ethernet synthetic frame loss measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.

#### Required Privilege Level

trace and maintenance

#### RELATED DOCUMENTATION

[clear oam ethernet connectivity-fault-management synthetic-loss-measurement](#) | 1446

[monitor ethernet synthetic-loss-measurement](#) | 1463

[show oam ethernet connectivity-fault-management loss-statistics](#) | 1896

#### List of Sample Output

[monitor ethernet synthetic-loss-measurement on page 1466](#)

#### Output Fields

The **monitor ethernet synthetic-loss-measurement** command displays different output at the CLI, depending on when you start a synthetic frame loss measurement:

- [Table 126 on page 1466](#) lists the run-time output fields for the **monitor ethernet synthetic-loss-measurement** command.

Output fields are listed in the approximate order in which they appear.

Table 126: monitor ethernet synthetic-loss-measurement Output Fields

| Output Field Name                        | Output Field Description                                                                                                                                                                                                                                                                                                          |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ETH-SLM request to</b>                | Unicast MAC address of the remote peer MEP.                                                                                                                                                                                                                                                                                       |
| <b>Interface</b>                         | Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.                                                                                                                                                                                                                                    |
| <b>SLM packets sent</b>                  | Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.                                                                                                                                                                                                   |
| <b>SLR packets received</b>              | Total number of synthetic loss reply (SLR) PDU frames received by the source MEP from the remote MEP during this measurement session.                                                                                                                                                                                             |
| <b>Local TXFC1 value</b>                 | Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.   |
| <b>Local RXFC1 value</b>                 | Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet. |
| <b>Last Received SLR frame TXFCf(tc)</b> | Value of the local counter TxFCI at the time of SLM frame transmission.                                                                                                                                                                                                                                                           |
| <b>Last Received SLR frame TXFCb(t)</b>  | Value of the local counter RxFCI at the time of SLR frame transmission.                                                                                                                                                                                                                                                           |
| <b>Frame loss (near-end)</b>             | Count of frame loss associated with ingress data frames.                                                                                                                                                                                                                                                                          |
| <b>Frame loss (far-end)</b>              | Count of frame loss associated with egress data frames.                                                                                                                                                                                                                                                                           |

## Sample Output

monitor ethernet synthetic-loss-measurement

```
user@host> monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```

ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0

Synthetic Loss measurement statistics:

SLM packets sent : 100

SLR packets received : 100

Accumulated SLM statistics:

Local TXFC1 value : 100

Local RXFC1 value : 100

Last Received SLR frame TXFCf(tc) : 100

Last Received SLR frame TXFCb(tc) : 100

SLM Frame Loss:

Frame Loss (far-end) : 0 (0.00 %)

Frame Loss (near-end) : 0 (0.00 %)

## monitor ethernet synthetic-loss-measurement

### Syntax

```
monitor ethernet synthetic-loss-measurement
maintenance-domain md-name
maintenance-association ma-name
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<local-mep mep-id>
<wait interval-milliseconds>
<priority 802.1p value>
<size>
<xml>
```

### Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.

**local-mep** option introduced in Junos OS Release 15.1

### Description

(MX Series routers) Start an ITU-T Y.1731 Ethernet synthetic loss measurement (ETH-SLM) session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. ETH-SLM (also called synthetic frame loss measurement) statistics are stored at one of the MEPs for later retrieval.

**NOTE:** If you attempt to monitor delays to a nonexistent MAC address, you must press Ctrl +c to explicitly quit the **monitor ethernet synthetic-loss-measurement** command and return to the CLI command prompt.

To start an Ethernet synthetic frame loss measurement session, the router initiates an exchange of frames carrying synthetic frame loss measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet synthetic frame loss measurement PDU frames exchanged to measure frame delay times—are displayed as the run-time output of the **monitor ethernet synthetic-loss-measurement** command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet synthetic frame loss measurement statistics, described below, are measured and stored at only one of the MEPs:

**Frame delay**—The difference, in microseconds, between the time a frame is sent and when it is received.

**Frame delay variation**—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

### Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**mep *remote-mep-id***—Numeric identifier of the peer MEP with which to perform Ethernet synthetic frame loss measurement. The discovered MAC address of the peer MEP is used. The range of values is from 1 through 8191.

***remote-mac-address***—Unicast MAC address of the peer MEP with which to perform Ethernet synthetic frame loss measurement. 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**. Multicast MAC addresses are not supported.

**count *frame-count***—(Optional) Number of frames to send to the specified peer MEP. The range of values is from 1 through 65,535 frames. The default value is 10 frames.

**local-mep *mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

**wait *interval-milliseconds***—(Optional) Number of milliseconds to wait between sending frames. You must specify this value in multiples of 100 milliseconds. The range of values is from 100 through 50000 milliseconds. The default value is 100 milliseconds.

**priority *802.1p value***—(Optional) Priority of the ETH-SLM request frame supported. The range of values is from 0 through 7. The default value is zero.

**size**—(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

**xml**—(Optional) Allow the output of the command to be displayed in XML format for ETH-SLM. Note that the only way to get output in XML format is to use the **xml** argument. The **display xml** command does not work.

### Additional Information

To display the frame counts collected at a MEP as the result of this command, use the following commands as described in the [CLI Explorer](#):

- **show oam ethernet connectivity-fault-management interfaces detail**
- **show oam ethernet connectivity-fault-management mep-database**
- **show oam ethernet connectivity-fault-management mep-statistics**

To display the statistics collected at a MEP as the result of this command, use the following commands as described in the [CLI Explorer](#).

- **show oam ethernet connectivity-fault-management synthetic-loss-measurement**
- **show oam ethernet connectivity-fault-management mep-statistics**

To clear both the frame counts and the statistics collected for MEPs, use the **clear oam ethernet connectivity-fault-management statistics** command, described in the [CLI Explorer](#).

For a complete description of Ethernet synthetic frame loss measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.

#### Required Privilege Level

trace and maintenance

#### RELATED DOCUMENTATION

[clear oam ethernet connectivity-fault-management synthetic-loss-measurement](#) | 1446

[monitor ethernet synthetic-loss-measurement](#) | 1463

[show oam ethernet connectivity-fault-management loss-statistics](#) | 1896

#### List of Sample Output

[monitor ethernet synthetic-loss-measurement on page 1471](#)

#### Output Fields

The **monitor ethernet synthetic-loss-measurement** command displays different output at the CLI, depending on when you start a synthetic frame loss measurement:

- [Table 126 on page 1466](#) lists the run-time output fields for the **monitor ethernet synthetic-loss-measurement** command.

Output fields are listed in the approximate order in which they appear.

Table 127: monitor ethernet synthetic-loss-measurement Output Fields

| Output Field Name                        | Output Field Description                                                                                                                                                                                                                                                                                                          |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ETH-SLM request to</b>                | Unicast MAC address of the remote peer MEP.                                                                                                                                                                                                                                                                                       |
| <b>Interface</b>                         | Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.                                                                                                                                                                                                                                    |
| <b>SLM packets sent</b>                  | Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.                                                                                                                                                                                                   |
| <b>SLR packets received</b>              | Total number of synthetic loss reply (SLR) PDU frames received by the source MEP from the remote MEP during this measurement session.                                                                                                                                                                                             |
| <b>Local TXFC1 value</b>                 | Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.   |
| <b>Local RXFC1 value</b>                 | Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet. |
| <b>Last Received SLR frame TXFCf(tc)</b> | Value of the local counter TxFCI at the time of SLM frame transmission.                                                                                                                                                                                                                                                           |
| <b>Last Received SLR frame TXFCb(t)</b>  | Value of the local counter RxFCI at the time of SLR frame transmission.                                                                                                                                                                                                                                                           |
| <b>Frame loss (near-end)</b>             | Count of frame loss associated with ingress data frames.                                                                                                                                                                                                                                                                          |
| <b>Frame loss (far-end)</b>              | Count of frame loss associated with egress data frames.                                                                                                                                                                                                                                                                           |

## Sample Output

monitor ethernet synthetic-loss-measurement

```
user@host> monitor ethernet synthetic-loss-measurement 00:05:85:73:39:4a maintenance-domain
md6 maintenance-association ma6 count 10
```



ETH-SLM request to 00:05:85:73:39:4a, interface ge-1/0/0.0

Synthetic Loss measurement statistics:

SLM packets sent : 100

SLR packets received : 100

Accumulated SLM statistics:

Local TXFC1 value : 100

Local RXFC1 value : 100

Last Received SLR frame TXFCf(tc) : 100

Last Received SLR frame TXFCb(tc) : 100

SLM Frame Loss:

Frame Loss (far-end) : 0 (0.00 %)

Frame Loss (near-end) : 0 (0.00 %)

## prbs-test-start

### Syntax

```
test interfaces ifd-name prbs-test-start pattern-type type direction 0/1 flip 0/1
```

### Release Information

Command introduced in Junos OS Release 19.2R1 for MX10003 and MX204 router.

### 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** 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 910.

### 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        |

| Pattern Type | Pattern Name                                                                |
|--------------|-----------------------------------------------------------------------------|
| 11           | PRBS11                                                                      |
| 15           | PRBS15                                                                      |
| 23           | PRBS23                                                                      |
| 31           | PRBS31                                                                      |
|              | <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

#### RELATED DOCUMENTATION

[prbs-test-stop | 1475](#)

[show interfaces prbs-stats | 1842](#)

[clear interfaces statistics](#)

[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test | 910](#)

## prbs-test-stop

### Syntax

```
test interfaces ifd-name prbs-test-stop direction (0|1)
```

### Release Information

Command introduced in Junos OS Release 19.2R1 for MX10003 and MX204 routers.

### Description

Use this command to stop collecting the Pseudo Random Binary Sequence (PRBS) statistics that is initiated using [prbs-test-start](#) 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 910.

### Options

***ifd-name***—Name of the interface.

***direction*** —Direction to transmit or receive PRBS pattern.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[prbs-test-start](#) | [1473](#)

[show interfaces prbs-stats](#) | [1842](#)

[clear interfaces statistics](#)

[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#) | [910](#)

### List of Sample Output

[Stopping PRBS test statistics collection on page 1475](#)

## Sample Output

### Stopping PRBS test statistics collection

```
user@host> test interface et-0/1/2 prbs-test-stop direction 1
```

## request interface link-degrade-recover

### Syntax

```
request interface link-degrade-recover interfaces-name
```

### Release Information

Command introduced in Junos OS Release 15.1.

### 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

### RELATED DOCUMENTATION

[Link Degrade Monitoring Overview | 843](#)

[link-degrade-monitor | 1227](#)

[thresholds | 1384](#)

[recovery | 1335](#)

### List of Sample Output

[Manual recovery on page 1477](#)

[Interface status when link degrade is enabled on page 1477](#)

[Interface status when the defect is active on page 1478](#)

### 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 Degradate :
    Link Monitoring          : Enable
    Link Degradate Set Threshold : 1E-8
    Link Degradate Clear Threshold : 1E-11
    Link Degradate War Set Threshold : 1E-9
    Link Degradate War Clear Threshold : 1E-10
    Estimated BER            : <= 1E-16
    Link-degrade event       : Seconds          Count
  State
    0                          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: Sendbroadcast-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 | Seconds |
|----------------|---------|
|----------------|---------|

|            |   |
|------------|---|
| Bit errors | 0 |
|------------|---|

|                |   |
|----------------|---|
| Errored blocks | 0 |
|----------------|---|

Link Degradate :

|                 |          |
|-----------------|----------|
| Link Monitoring | : Enable |
|-----------------|----------|

|                              |        |
|------------------------------|--------|
| Link Degradate Set Threshold | : 1E-8 |
|------------------------------|--------|

|                                |         |
|--------------------------------|---------|
| Link Degradate 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

```



## request interface mc-ae switchover (Multichassis Link Aggregation)

### Syntax

```
request interface mc-ae switchover
<immediate> mcae-id mcae-id;
mcae-id mcae-id;
```

### Release Information

Command introduced in Junos OS Release 13.3.

### 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

### RELATED DOCUMENTATION

*Configuring Multichassis Link Aggregation on MX Series Routers*

*Configuring Manual and Automatic Link Switchover for MC-LAG Interfaces on MX Series Routers*

### List of Sample Output

[request interface mc-ae switchover immediate mcae-id on page 1481](#)

[request interface mc-ae switchover mcae-id on page 1481](#)

### 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
```

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

### Syntax

```
request interface (revert | switchover) aex
```

### Release Information

Command introduced in Junos OS Release 8.3.

### 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

**revert**—Restores egress traffic processing to the primary link.

**switchover**—Transfers egress traffic processing to the secondary (backup) link.

**aex**—Aggregated Ethernet logical interface number: 0 through 15.

### Required Privilege Level

view

### List of Sample Output

[request interface revert on page 1483](#)

### 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
```

## request lacp link-switchover

### Syntax

```
request lacp link-switchover aex
```

### Release Information

Command introduced in Junos OS Release 9.3.

### 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

### List of Sample Output

[request lacp link-switchover aeX on page 1484](#)

### 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
```

## show chassis hardware

### List of Syntax

[Syntax on page 1485](#)

[Syntax \(EX Series, MX104, MX204, MX2010, MX2020, MX10003, MX10008, and MX2008 Universal Routing Platforms\) on page 1485](#)

[Syntax \(TX Matrix Router\) on page 1485](#)

[Syntax \(TX Matrix Plus Router\) on page 1485](#)

[Syntax \(MX Series Routers\) on page 1486](#)

[Syntax \(QFX Series\) on page 1486](#)

### 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>
```

## Release Information

Command introduced before Junos OS Release 7.4.

**models** option introduced in Junos OS Release 8.2.

Command introduced in Junos OS Release 9.0 for EX Series switches.

**sfc** option introduced in Junos OS Release 9.6 for the TX Matrix Plus router.

Command introduced in Junos OS Release 11.1 for QFX Series.

Command introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 Universal Routing Platforms.

Information for **disk** and **usb** introduced in Junos OS Release 15.1X53-D60 for QFX10002, QFX10008, and QFX10016 switches.

Command introduced in Junos OS Release 15.1X54-D20 for ACX5048 and ACX5096 Routers.

Command introduced in Junos OS Release 17.2 for MX2008 Universal Routing Platforms and PTX10008 Routers.

Command introduced in Junos OS Release 17.3 for MX10003 Universal Routing Platforms and MX150 Router Appliance.

Command introduced in Junos OS Release 17.4 for MX204 Routers.

Command introduced in Junos OS Release 18.1R1 for EX9251 Switches.

Command introduced in Junos OS Release 18.2R1 for MX10008 Routers and EX9253 Switches.

**NOTE:** Routers and routing platforms use the basic syntax, unless otherwise listed. For example, the EX Series has an additional satellite parameter available.

## 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

**none**—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.

**clei-models**—(Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).

**detail**—(Optional) Include RAM and disk information in output.

**extensive**—(Optional) Display ID EEPROM information.

**all-members**—(MX Series routers only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.



**interconnect-device *name***—(QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.

**lcc *number***—(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.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 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.

**local**—(MX Series routers only) (Optional) Display hardware-specific information for the local Virtual Chassis members.

**member *member-id***—(MX Series routers and EX Series switches) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value 0 or 1.

**models**—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

**node-device *name***—(QFabric systems only) (Optional) Display hardware-specific information for the Node device.

**satellite [*slot-id slot-id* | *device-alias alias-name*]**—(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.

**scc**—(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

### Additional Information

The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines, as shown in [Table 128 on page 1489](#).

Table 128: 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

#### RELATED DOCUMENTATION

| [show chassis power](#)

#### List of Sample Output

[show chassis hardware \(MX10008 Router\) on page 1494](#)  
[show chassis hardware clei-models \(PTX10016 Routers\) on page 1495](#)  
[show chassis hardware detail \(EX9251 Switch\) on page 1497](#)  
[show chassis hardware extensive \(T640 Router\) on page 1497](#)  
[show chassis hardware interconnect-device \(QFabric Systems\) on page 1498](#)  
[show chassis hardware lcc \(TX Matrix Router\) on page 1499](#)  
[show chassis hardware models \(MX2010 Router\) on page 1500](#)  
[show chassis hardware node-device \(QFabric Systems\) on page 1501](#)  
[show chassis hardware scc \(TX Matrix Router\) on page 1501](#)  
[show chassis hardware sfc \(TX Matrix Plus Router\) on page 1501](#)

#### Output Fields

[Table 129 on page 1490](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 129: show chassis hardware Output Fields

| Field Name | Field Description | Level of Output |
|------------|-------------------|-----------------|
| Item       |                   | All levels      |

Table 129: show chassis hardware Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <p>Chassis component:</p> <ul style="list-style-type: none"> <li>• (EX Series switches)—Information about the chassis, Routing Engine (SRE and Routing Engine modules in EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for EX8200 switches.</li> <li>• (MX Series routers and EX Series switches)—Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs). MX80 routers have a single Routing Engine and a built-in Packet Forwarding Engine that attaches directly to MICs. The Packet Forwarding Engine has two “pseudo” FPCs (FPC 0 and FPC1). MX80 routers also have a Forwarding Engine Board (FEB). MX104 routers have a built-in Packet forwarding Engine and a Forwarding Engine Board (FEB). The Packet Forwarding Engine of the MX104 router has three “pseudo” FPCs (FPC0, FPC1, and FPC2).</li> <li>• (M Series routers, except for the M320 router)—Information about the backplane; power supplies; fan trays; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); SCB, SSB, SFM, or FEB; MCS and PCG (for the M160 router only); each FPC and PIC; and each fan, blower, and impeller.</li> <li>• (M120, M320, and T Series routers)—Information about the backplane, power supplies, fan trays, midplane, FPM (craft interface), CIP, PEM, SCG, CB, FPC, PIC, SFP, SPMB, and SIB.</li> <li>• (QFX Series)—Information about the chassis, Pseudo CB, Routing Engine, power supplies, fan trays, Interconnect devices, and Node devices. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs).</li> <li>• (PTX Series)—Information about the chassis, midplane, craft interface (FPM), power distribution units (PDUs) and Power Supply Modules (PSMs), Centralized Clock Generators (CCGs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Flexible PIC Concentrators (FPCs), PICs, Switch Interface Boards (SIBs), and fan trays (vertical and horizontal).</li> <li>• (MX2010, MX2020, and MX2008 routers)—Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing</li> </ul> |                 |

Table 129: show chassis hardware Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output  |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                               | Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays.<br><ul style="list-style-type: none"> <li>(vMX routers)—Information about the chassis, midplane, Routing Engines, and Control Boards (CBs). Also displays information about Flexible PIC Concentrators (FPCs) and associated Modular Interface Cards (MICs) and Physical Interface Cards (PICs).</li> </ul> |                  |
| <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> |
| <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> |
| <b>Description</b>            | Brief description of the hardware item:<br><ul style="list-style-type: none"> <li>Type of power supply.</li> <li>Type of PIC. If the PIC type is not supported on the current software release, the output states <b>Hardware Not Supported</b>.</li> </ul>                                                                                                                                                                                                                                        | All levels       |

Table 129: show chassis hardware Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <li>• Type of FPC: <b>FPC Type 1</b>, <b>FPC Type 2</b>, <b>FPC Type 3</b>, <b>FPC Type 4</b> , or <b>FPC TypeOC192</b>.</li> </ul> <p>On EX Series switches, a brief description of the FPC.</p> <p>The following list shows the PIM abbreviation in the output and the corresponding PIM name.</p> <ul style="list-style-type: none"> <li>• <b>2x FE</b>—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM</li> <li>• <b>4x FE</b>—4-port Fast Ethernet ePIM</li> <li>• <b>1x GE Copper</b>—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port)</li> <li>• <b>1x GE SFP</b>—SFP Gigabit Ethernet ePIM (one fiber port)</li> <li>• <b>2x Serial</b>—Dual-port serial PIM</li> <li>• <b>2x T1</b>—Dual-port T1 PIM</li> <li>• <b>2x E1</b>—Dual-port E1 PIM</li> <li>• <b>2x CT1E1</b>—Dual-port channelized T1/E1 PIM</li> <li>• <b>1x T3</b>—T3 PIM (one port)</li> <li>• <b>1x E3</b>—E3 PIM (one port)</li> <li>• <b>4x BRI S/T</b>—4-port ISDN BRI S/T PIM</li> <li>• <b>4x BRI U</b>—4-port ISDN BRI U PIM</li> <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, 4x FX0, 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> </ul> |                 |

Table 129: show chassis hardware Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <ul style="list-style-type: none"> <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</b>, <b>SX</b>, <b>LH</b>, or <b>T</b>.</li> <li>• LCD description for EX Series switches (except EX2200 switches).</li> <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                REV 27   750-054097   ACPD4307      JNP10008 [MX10008]
Midplane                REV 02   750-079563   CAFF4580      Control Board
Routing Engine 0                BUILTIN     BUILTIN       RE X10 LT
Routing Engine 1                BUILTIN     BUILTIN       RE X10
CB 0                   REV 02   750-079563   CAFF4580      Control Board
CB 1                   REV 04   750-079563   CAGL8034      Control Board
..

```

```

...
..
4
FPC 3          REV 04  750-084779  CAKR7019          JNP10K-LC2101
  CPU          REV 05  750-073391  CAKJ2854          LC 2101 PMB
  PIC 0                BUILTIN    BUILTIN           4xQSFP28 SYNCE
    Xcvr 0        REV 01  740-058734  1ACQ104300K       QSFP-100GBASE-SR4
  PIC 1                BUILTIN    BUILTIN           4xQSFP28 SYNCE
    Xcvr 0        REV 01  740-061405  1ACQ12110AN       QSFP-100GBASE-SR4
  PIC 2                BUILTIN    BUILTIN           4xQSFP28 SYNCE
    Xcvr 0        REV 01  740-046565  QG1105B2          QSFP+-40G-SR4
  PIC 3                BUILTIN    BUILTIN           4xQSFP28 SYNCE
    Xcvr 0        REV 01  740-045627  QH08036X          40GBASE eSR4
  PIC 4                BUILTIN    BUILTIN           4xQSFP28 SYNCE
    Xcvr 0        REV 01  740-067443  XWR0RY7           QSFP+-40G-SR4
    Xcvr 1        REV 01  740-067443  XWR0RYH           QSFP+-40G-SR4
    Xcvr 2        REV 01  740-067443  XWR0RYP           QSFP+-40G-SR4
    Xcvr 3        REV 01  740-067443  XWS028S           QSFP+-40G-SR4
  PIC 5                BUILTIN    BUILTIN           4xQSFP28 SYNCE
    Xcvr 3        REV 01  740-058734  1ACQ113406C       QSFP-100GBASE-SR4
FPD Board      REV 07  711-054687  ACPD6799          Front Panel Display
PEM 0          REV 02  740-049388  1EDL62102N9       Power Supply AC
PEM 1          REV 02  740-049388  1EDL60300KX       Power Supply AC
PEM 2          REV 02  740-049388  1EDL60300DL       Power Supply AC
PEM 3          REV 02  740-049388  1EDL61701BT       Power Supply AC
PEM 4          REV 02  740-049388  1EDL62102P7       Power Supply AC
PEM 5          REV 02  740-049388  1EDL62102PP       Power Supply AC
FTC 0          REV 14  750-050108  ACPE4038          Fan Controller 8
FTC 1          REV 14  750-050108  ACPE4032          Fan Controller 8
Fan Tray 0     REV 09  760-054372  ACPD6799          Fan Tray 8
Fan Tray 1     REV 09  760-054372  ACNZ3584          Fan Tray 8
SFB 0          REV 24  750-050058  ACPD4587          Switch Fabric (SIB) 8
SFB 1          REV 24  750-050058  ACNZ0635          Switch Fabric (SIB) 8
SFB 2          REV 24  750-050058  ACPD4908          Switch Fabric (SIB) 8
SFB 3          REV 24  750-050058  ACNZ0617          Switch Fabric (SIB) 8
SFB 4          REV 24  750-050058  ACNZ0527          Switch Fabric (SIB) 8
SFB 5          REV 23  750-050058  ACNX6980          Switch Fabric (SIB) 8

```

### show chassis hardware clei-models (PTX10016 Routers)

```
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          | REV 13 | 750-068822 | CMUIAM9BAC | QFX10000-36Q    |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 6          | REV 41 | 750-071976 | CMUIANABAB | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 7          | REV 35 | 750-071976 | CMUIANABAA | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 8          | REV 35 | 750-071976 | CMUIANABAA | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 9          | REV 41 | 750-071976 | CMUIANABAB | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 10         | REV 35 | 750-071976 | CMUIANABAA | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 11         | REV 35 | 750-071976 | CMUIANABAA | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 13         | REV 41 | 750-071976 | CMUIANABAB | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| FPC 15         | REV 37 | 750-071976 | CMUIANABAA | JNP10K-LC1101   |
| PIC 0          |        | BUILTIN    |            |                 |
| 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        4XQSFP28 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    13T511102684 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 ff ff ff ff ff ff ff ff ff ff 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 ff ff ff ff ff ff ff ff ff ff ff ff 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  |             |               | QFX_olive    |
| 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          |         |             | 65751         | T640                    |
| 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  | qbl3984       | 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    |
| SFP 0            | REV 01  | 740-009898  | USC202U709    | XENPAK-LR               |
| PIC 1            | REV 03  | 750-003336  | HJ9954        | 4x OC-48 SONET, SMSR    |
| PIC 2            | REV 01  | 750-004535  | HC0235        | 1x OC-192 SM SR1        |
| PIC 3            | REV 07  | 750-007141  | HX1699        | 10x 1GE(LAN), 1000 BASE |
| SFP 0            | REV 01  | 740-007326  | 2441042       | SFP-SX                  |
| SFP 1            | REV 01  | 740-007326  | 2441027       | SFP-SX                  |

|         |        |            |        |                     |
|---------|--------|------------|--------|---------------------|
| MMB 0   | REV 03 | 710-010171 | HV2365 | MMB-5M3-288mbit     |
| MMB 1   | REV 03 | 710-010171 | HZ3888 | MMB-5M3-288mbit     |
| SPMB 0  | REV 09 | 710-003229 | HW5245 | T Series Switch CPU |
| SIB 3   | REV 07 | 710-005781 | HR5927 | SIB-L8-F16          |
| B Board | REV 06 | 710-005782 | HR5971 | SIB-L8-F16 (B)      |
| SIB 4   | REV 07 | 710-005781 | HR5903 | SIB-L8-F16          |
| B Board | REV 06 | 710-005782 | HZ5275 | SIB-L8-F16 (B)      |

### show chassis hardware models (MX2010 Router)

user@host > show chassis hardware models

#### Hardware inventory:

| Item             | Version | Part number | Serial number | FRU model number       |
|------------------|---------|-------------|---------------|------------------------|
| FPM Board        | REV 06  | 711-032349  | ZX8744        | 711-032349             |
| PSM 4            | REV 0C  | 740-033727  | VK00254       | 0000000000000000000000 |
| PSM 5            | REV 0B  | 740-033727  | VG00015       | 0000000000000000000000 |
| PSM 6            | REV 0B  | 740-033727  | VH00097       | 0000000000000000000000 |
| PSM 7            | REV 0C  | 740-033727  | VJ00151       | 0000000000000000000000 |
| PSM 8            | REV 0C  | 740-033727  | VJ00149       | 0000000000000000000000 |
| PDM 0            | REV 0B  | 740-038109  | WA00008       |                        |
| PDM 1            | REV 0B  | 740-038109  | WA00014       |                        |
| Routing Engine 0 | REV 02  | 740-041821  | 9009094134    | RE-S-1800X4-16G-S      |
| Routing Engine 1 | REV 02  | 740-041821  | 9009094141    | RE-S-1800X4-16G-S      |
| CB 0             | REV 08  | 750-040257  | CAAB3491      | 750-040257             |
| CB 1             | REV 08  | 750-040257  | CAAB3489      | 750-040257             |
| SFB 0            | REV 06  | 711-032385  | ZV1828        | 711-032385             |
| SFB 1            | REV 07  | 711-032385  | ZZ2568        | 711-032385             |
| SFB 2            | REV 07  | 711-032385  | ZZ2563        | 711-032385             |
| SFB 3            | REV 07  | 711-032385  | ZZ2564        | 711-032385             |
| SFB 4            | REV 07  | 711-032385  | ZZ2580        | 711-032385             |
| SFB 5            | REV 07  | 711-032385  | ZZ2579        | 711-0323856            |
| SFB 6            | REV 07  | 711-032385  | CAAB4882      | 711-044170             |
| SFB 7            | REV 07  | 711-032385  | CAAB4898      | 711-044170             |
| FPC 0            | REV 33  | 750-028467  | CAAB1919      | MPC-3D-16XGE-SFPP      |
| FPC 1            | REV 21  | 750-033205  | ZG5027        | MX-MPC3-3D             |
| 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                  |
| PIC 0            |         | BUILTIN    | BUILTIN                  |
| Xcvr 8           | REV 01  | 740-030658 | AD0946A028B              |
|                  |         |            | FPC CPU                  |
|                  |         |            | 48x 10G-SFP+             |
|                  |         |            | 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                               TX Matrix
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          |         |             | JN112F007AHB  | TXP                |
| Midplane         | 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  |



## show chassis pic

### List of Syntax

[Syntax on page 1504](#)

[Syntax \(TX Matrix and TX Matrix Plus Routers\) on page 1504](#)

[Syntax \(MX Series Routers and EX Series Switches\) on page 1504](#)

[Syntax \(PTX Series Packet Transport Router and MX240, MX480, MX960, MX2010, and MX2020 Routers\) on page 1504](#)

[Syntax \(QFX Series\) on page 1504](#)

[Syntax \(ACX5048 and ACX5096 Routers\) on page 1504](#)

[Syntax \(ACX500 Routers\) on page 1505](#)

### 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
<icc 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)
```

### Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for QFX Series.

Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Command introduced in Junos OS Release 12.3 for MX2010 and MX2020 3D Universal Edge Routers.

Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers and MX104 3D Universal Edge Routers.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**transport** option introduced in Junos OS Release 16.1R1 for MX Series Routers.

Command introduced in Junos OS Release 17.2 for MX2008 3D Universal Edge Routers and PTX10008 Routers.

Command introduced in Junos OS Release 17.3 for MX10003 3D Universal Edge Routers and MX150 Router Appliance.

Command introduced in Junos OS Release 17.4 for MX204 3D Universal Edge Routers.

### 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.

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

**interconnect-device name**—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

**lcc number**—(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.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 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.

**local**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

**pic-slot *slot-number***—Display information about the PIC in this particular PIC slot. For routers, replace *slot-number* with a value from 0 through 3. For EX3200 and EX4200 switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

**transport**—Display PIC information for optical transport network.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[request chassis pic](#)

[show chassis hardware](#) | 1485

[100-Gigabit Ethernet Type 4 PIC with CFP Overview](#) | 247

### List of Sample Output

[show chassis pic fpc-slot pic-slot on page 1512](#)

[show chassis pic fpc-slot pic-slot \(PIC Offline\) on page 1512](#)

[show chassis pic fpc-slot pic-slot \(FPC Offline\) on page 1513](#)

[show chassis pic fpc-slot pic-slot \(FPC Not Present\) on page 1513](#)

[show chassis pic fpc-slot pic-slot \(PIC Not Present\) on page 1513](#)

[show chassis pic fpc-slot 3 pic-slot 0 \(M120 Router\) on page 1513](#)

[show chassis pic fpc-slot pic-slot \(MX150\) on page 1513](#)  
[show chassis pic fpc-slot pic-slot \(MX960 Router with Bidirectional Optics\) on page 1514](#)  
[show chassis pic fpc-slot pic-slot \(MX480 Router with 100-Gigabit Ethernet MIC\) on page 1514](#)  
[show chassis pic fpc-slot pic-slot \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 1515](#)  
[show chassis pic fpc-slot pic-slot \(MX960 Router with MPC5EQ\) on page 1515](#)  
[show chassis pic fpc-slot pic-slot \(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC\) on page 1516](#)  
[show chassis pic fpc-slot pic-slot \(MX10003 Routers\) on page 1516](#)  
[show chassis pic fpc-slot pic-slot \(PTX1000 and PTX10000\) on page 1517](#)  
[show chassis pic fpc-slot pic-slot \(PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC on page 1519](#)  
[show chassis pic fpc-slot pic-slot \(MX480 Router with MPC4E\) on page 1520](#)  
[show chassis pic fpc-slot pic-slot \(MX480 router with OTN Interface\) on page 1520](#)  
[show chassis pic fpc-slot pic-slot \(MX2010 Router with OTN Interfaces\) on page 1521](#)  
[show chassis pic fpc-slot pic-slot \(MX2010 Router\) on page 1521](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Router\) on page 1521](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC5EQ and MPC6E\) on page 1522](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC6E and OTN MIC\) on page 1522](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Router with MPC4E\) on page 1523](#)  
[show chassis pic fpc-slot pic-slot \(MX2010 Router\) on page 1524](#)  
[show chassis pic fpc-slot pic-slot \(T1600 Router with 100-Gigabit Ethernet PIC\) on page 1524](#)  
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Router\) on page 1524](#)  
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Plus Router\) on page 1525](#)  
[show chassis pic fpc-slot pic-slot \(Next-Generation SONET/SDH SFP\) on page 1525](#)  
[show chassis pic fpc-slot pic-slot \(12-Port T1/E1\) on page 1525](#)  
[show chassis pic fpc-slot 0 pic-slot 1 \(4x CHOC3 SONET CE SFP\) on page 1526](#)  
[show chassis pic fpc-slot 0 pic-slot 0 \(SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 1526](#)  
[show chassis pic fpc-slot 3 pic-slot 0 \(8-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 1527](#)  
[show chassis pic fpc-slot 5 pic-slot 0 \(4-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 1527](#)  
[show chassis pic fpc-slot 1 pic-slot 0 \(1-port OC192/STM64 MIC with XFP\) on page 1528](#)  
[show chassis pic fpc-slot 1 pic-slot 2 \(8-port DS3/E3 MIC\) on page 1528](#)  
[show chassis pic fpc-slot pic-slot \(OTN\) on page 1528](#)  
[show chassis pic fpc-slot pic-slot \(QFX3500 Switch\) on page 1529](#)  
[show chassis pic fpc-slot pic-slot \(QFX5100 Switches and OCX Series \) on page 1529](#)  
[show chassis pic interconnect-device fpc-slot pic-slot \(QFabric Systems\) on page 1529](#)  
[show chassis pic node-device fpc-slot pic-slot \(QFabric System\) on page 1529](#)  
[show chassis pic fpc-slot 0 pic-slot 1 \(ACX2000 Universal Access Router\) on page 1531](#)  
[show chassis pic FPC-slot 1 PIC-slot 0 \(MX Routers with Media Services Blade \[MSB\]\) on page 1531](#)  
[show chassis pic FPC slot 1, PIC slot 2 \(MX Routers with Media Services Blade \[MSB\]\) on page 1531](#)  
[show chassis pic transport fpc-slot pic-slot \(PTX Series Packet Transport Routers\) on page 1531](#)

[show chassis pic transport fpc-slot pic-slot \(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC\) on page 1531](#)

[show chassis pic fpc-slot 0 pic-slot 0 \(ACX5096 Router\) on page 1532](#)

[show chassis pic fpc-slot 0 pic-slot 0 \(ACX5048 Router\) on page 1537](#)

[show chassis pic fpc-slot 0 pic-slot 0 \(ACX500 Router\) on page 1537](#)

[show chassis pic fpc-slot 0 pic-slot 1 \(ACX500 Router\) on page 1537](#)

[show chassis pic transport fpc-slot pic-slot \(PTX Series Packet Transport Routers\) on page 1538](#)

[show chassis pic transport fpc-slot pic-slot \(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC\) on page 1538](#)

[show chassis pic fpc-slot 7 pic-slot 1 \(MX960 Router MPC10E-15C-MRATE Line Card\) on page 1538](#)

## Output Fields

[Table 130 on page 1509](#) lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 130: show chassis pic Output Fields

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                           |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Type                    | <p>PIC type.</p> <p><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>.</p> |
| Account Layer2 Overhead | (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.                                                                                                                                                                |
| ASIC type               | Type of ASIC on the PIC.                                                                                                                                                                                                                                                                                                    |

Table 130: 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>                                                      | (Multiservices PICs only) Services package supported: <b>Layer-2</b> or <b>Layer-3</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <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> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)</b> | <p>Port-level information for the PIC.</p> <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>                                                                                                                                                       |

Table 130: show chassis pic Output Fields (*continued*)

| Field Name                                                               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>PIC Port Information<br/>(MX960 Router<br/>Bidirectional Optics )</b> | <p>Port-level information for the PIC.</p> <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> |
| <b>PIC Port Information<br/>(Next-Generation<br/>SONET/SDH SFP)</b>      | <p>Port-level information for the next-generation SONET/SDH SFP PIC.</p> <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<br/>(MX104 router)</b>                           | <p>Port-level information for the PIC.</p> <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>                                                                                                                                                                                                                                                                               |



Table 130: show chassis pic Output Fields (*continued*)

| Field Name             | Field Description                                                                                                                                                                                                            |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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> |
| Multirate Mode         | Rate-selectability status for the MIC: <b>Enabled</b> or <b>Disabled</b> .                                                                                                                                                   |
| Channelization         | Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.                                                                                                                                                   |
| Administrative State   | Indicates the administrative state of the PIC. Possible values are: In Service (Default) and Out of Service.                                                                                                                 |
| Operational State      | 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:

|      |            | Fiber |               | Xcvr vendor  | Wave-  | Xcvr |
|------|------------|-------|---------------|--------------|--------|------|
| Port | Cable type | type  | Xcvr vendor   | part number  | length |      |
| 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 type | Xcvr vendor      | Xcvr vendor part number | Wavelength |
|------|---------------------|------------|------------------|-------------------------|------------|
| 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         | OCF              | TRXBG1LXDBVM2-JW        | 1490 nm    |
| 4    | SFP-1000BASE-BX10-D | SM         | OCF              | 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         | OCF              | TRXBG1LXDBBMH-J1        | 1310 nm    |
| 8    | SFP-1000BASE-BX10-U | SM         | OCF              | TRXBG1LXDBBMH-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:

```

      Fiber                Xcvr vendor      Wave-    Xcvr
      Port Cable type          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:

      Fiber                Xcvr vendor      Wave-      Xcvr
      Port Cable type      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:

      Fiber                Xcvr vendor      Wave-      Xcvr
      Port Cable type      type  Xcvr vendor      part number  length
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:

|          |              | Fiber | Xcvr vendor    | Wave-           | Xcvr        |
|----------|--------------|-------|----------------|-----------------|-------------|
| Port     | Cable type   | type  | Xcvr vendor    | part number     | length      |
| 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:

|                 |            | Fiber |                 | Xcvr vendor  | Wave-      | Xcvr |
|-----------------|------------|-------|-----------------|--------------|------------|------|
| Port            | Cable type | type  | Xcvr vendor     | part number  | length     |      |
| 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:

```

|          |            | Fiber |              | Xcvr vendor    | Wave-  | Xcvr |
|----------|------------|-------|--------------|----------------|--------|------|
| Port     | Cable type | type  | Xcvr vendor  | part number    | length |      |
| 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:

```

|          |            | Fiber |               | Xcvr vendor      | Wave-  | Xcvr |
|----------|------------|-------|---------------|------------------|--------|------|
| Port     | Cable 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:

      Fiber                Xcvr vendor      Wave-      Xcvr
      Port Cable 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: |            |                                         |                  |                |        |       |      |
|                       |            | Fiber                                   |                  | Xcvr vendor    |        | Wave- | Xcvr |
| Port                  | Cable type | type                                    | Xcvr vendor      | part number    | length |       |      |
| Firmware              |            |                                         |                  |                |        |       |      |
| 0                     | 10GBASE SR | MM                                      | SumitomoElectric | SPP5200SR-J6-M | 850 nm | 0.0   |      |
| 1                     | 10GBASE SR | MM                                      | SumitomoElectric | SPP5200SR-J6-M | 850 nm | 0.0   |      |
| 2                     | 10GBASE SR | MM                                      | SumitomoElectric | SPP5200SR-J6-M | 850 nm | 0.0   |      |
| 3                     | 10GBASE SR | MM                                      | SumitomoElectric | SPP5200SR-J6-M | 850 nm | 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 QSFPP          |             |                  |        |       |      |
| State                                |             | Online                |             |                  |        |       |      |
| PIC version                          |             | 0.0                   |             |                  |        |       |      |
| Uptime                               |             | 6 minutes, 31 seconds |             |                  |        |       |      |
| PIC port information:                |             |                       |             |                  |        |       |      |
|                                      |             | Fiber                 |             | Xcvr vendor      |        | Wave- | Xcvr |
| Port                                 | Cable type  | 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:

|          |            |            | Fiber            |                  | Xcvr vendor | Wave-  | Xcvr |
|----------|------------|------------|------------------|------------------|-------------|--------|------|
|          | Port       | Cable type | 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:

|          |            |            | Fiber            |              | Xcvr vendor | Wave-  | Xcvr |
|----------|------------|------------|------------------|--------------|-------------|--------|------|
|          | Port       | Cable type | 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      Xcvr vendor      Wavelength
  0      100GBASE 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    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    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      Xcvr vendor      Xcvr vendor      Wavelength
      type              type  Xcvr vendor      part number
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:

  Fiber      Xcvr vendor

```

| Port           | Cable type       | type | 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 type | Xcvr vendor   | part number      | Wavelength |
|------|------------------|------------|---------------|------------------|------------|
| 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 type | Xcvr vendor   | part number      | Wavelength |
|------|------------------|------------|---------------|------------------|------------|
| 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:
  Port  Cable type      Fiber 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-GE Builtin
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     |

[illegible]

**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:
```

|          |            |    | Fiber            |                  | Xcvr vendor | Wave-  | Xcvr |
|----------|------------|----|------------------|------------------|-------------|--------|------|
| Port     | Cable type |    | type             | Xcvr vendor      | part number | length |      |
| Firmware |            |    |                  |                  |             |        |      |
| 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     | 0.0 |

**show chassis pic fpc-slot 0 pic-slot 0 (ACX5048 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               1 day, 5 hours, 27 minutes, 25 seconds

```

```
PIC port information:
```

|          |             | Fiber |               | Xcvr vendor      | Wave-  | Xcvr |
|----------|-------------|-------|---------------|------------------|--------|------|
| Port     | Cable type  | type  | Xcvr vendor   | part number      | length |      |
| Firmware |             |       |               |                  |        |      |
| 0        | 10GBASE SR  | MM    | FINISAR CORP. | FTLX8571D3BCL-J1 | 850 nm | 0.0  |
| 10       | GIGE 1000SX | MM    | FINISAR CORP. | FTLF8519P3BNL-J1 | 850 nm | 0.0  |
| 14       | 10GBASE SR  | MM    | FINISAR CORP. | FTLX8571D3BNL-J1 | 850 nm | 0.0  |
| 20       | 10GBASE SR  | MM    | FINISAR CORP. | FTLX8571D3BCL-J1 | 850 nm | 0.0  |
| 30       | GIGE 1000SX | MM    | FINISAR CORP. | FTLF8519P2BNL-J1 | 850 nm | 0.0  |
| 41       | 10GBASE SR  | MM    | OPNEXT, INC.  | TRS2001EN-0014   | 850 nm | 0.0  |
| 46       | GIGE 1000SX | MM    | FINISAR CORP. | FTLF8519P2BNL-J1 | 850 nm | 0.0  |
| 64       | 10GBASE SR  | MM    | FINISAR CORP. | FTLX8571D3BNL-J1 | 850 nm | 0.0  |
| 78       | GIGE 1000SX | MM    | AVAGO         | AFBR-5715PZ-JU2  | 850 nm | 0.0  |
| 96       | 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 SR4 | MM    | AVAGO         | AFBR-79EQDZ-JU1  | 850 nm | 0.0  |

**show chassis pic fpc-slot 0 pic-slot 0 (ACX500 Router)**

```
user@host> show chassis pic fpc-slot 0 pic-slot 0
```

```
FPC slot 0, PIC slot 0 information:
```

```

Type                2x 1GE(LAN) SFP Builtin
State                Online
Uptime               17 hours, 54 minutes, 45 seconds

```

**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 Builtin
  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:
      Fiber                               Xcvr vendor      Wave-      Xcvr
      JNPR
  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 |

## show ethernet-switching redundancy-groups

### Syntax

```
show ethernet-switching redundancy-groups  
<redundancy-group-id [0 to 4294967294]>  
arp-statistics  
nd-statistics  
remote-macs
```

### Release Information

Command introduced in Junos OS Release 13.2.

Command introduced in Junos OS Release 15.1R1 for EX Series switches

### 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

**redundancy-group-id**—(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.

**arp-statistics**—(Optional) Count of ARP packets sent and received by the two MC-AE nodes.

**nd-statistics**—(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.

**remote-macs** —(Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.

### Required Privilege Level

view

## RELATED DOCUMENTATION

| [Configuring Multichassis Link Aggregation on EX Series Switches](#)

### List of Sample Output

[show ethernet-switching redundancy-groups arp-statistics on page 1542](#)

[show ethernet-switching redundancy-groups nd-statistics on page 1543](#)

[show ethernet-switching redundancy-groups remote-macs on page 1543](#)

[show ethernet-switching redundancy-groups group-id on page 1544](#)

## Output Fields

Output fields are listed in the approximate order in which they appear.

**Table 131: 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 132: show ethernet-switching 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 (MC-LAG) 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.                                  |

Table 132: show ethernet-switching redundancy-groups nd-statistics Output Fields (*continued*)

| Field Name                       | Field Description                                                               |
|----------------------------------|---------------------------------------------------------------------------------|
| 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 133: 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.                                                                                     |
| 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         | 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
```

## show interfaces (Adaptive Services)

### Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Release Information

Command introduced before Junos OS Release 7.4.

### 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/pic/port***.

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

**descriptions**—(Optional) Display interface description strings.

**media**—(Optional) Display media-specific information about network interfaces.

**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

### List of Sample Output

[show interfaces \(Adaptive Services\) on page 1549](#)

[show interfaces brief \(Adaptive Services\) on page 1550](#)

[show interfaces detail \(Adaptive Services\) on page 1550](#)

[show interfaces extensive \(Adaptive Services\) on page 1551](#)

### Output Fields

[Table 134 on page 1546](#) 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 134: 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 under <i>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.                                                                                  |
| <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>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>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.                                                                                                                            |
| <b>Hardware address</b>   | MAC address of the hardware.                                                                                                                       |

Table 134: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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 <b>Last flapped <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i></b> . For example, <b>Last flapped 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 MX Series routers and MS-MPCs, the Input bytes field might show 16 more bytes than the Output bytes field because of the accounting of 16 bytes of the Juniper Forwarding Module 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>                                                                                                                                                                                                                                          |
| <b>Input errors</b>            | <p>Input errors on the interface. The following paragraphs explain the counters whose meanings are obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the queue is saturated, this number increments once for every packet that is dropped by the 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>—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 support.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> |

Table 134: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Output errors</b>      | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning is not 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 normally increments quickly, increasing only when the cable is unplugged, the far-end system goes down and then up, or another problem occurs. If the number of carrier transitions increments slowly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the queue is saturated, this number increments once for every packet that is dropped by the ASIC.</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> |
| <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> | <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>Local statistics</b>   | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic occurs, the value in the output packet rate field might briefly exceed the peak cell rate. It takes less than 1 second for this counter to stabilize.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 134: Adaptive Services and Redundant Adaptive Services show interfaces Output Fields *(continued)*

| Field Name                    | Field Description                                                                                                                                                                                                         |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Transit statistics</b>     | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the rate field might briefly exceed the peak cell rate. It takes generally less than 1 second for the value to stabilize. |
| <b><i>protocol-family</i></b> | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address 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> refers to the inet.0.                                                                                                              |
| <b>Flags</b>                  | Information about the protocol family flags. Possible values are described in the “Family Flags” section of the <i>Common Output Fields Description</i> .                                                                 |
| <b>Addresses, Flags</b>       | Information about the address flags. Possible values are described in the “Addresses and Flags” section of the <i>Common Output Fields Description</i> .                                                                  |
| <b>Destination</b>            | IP address of the remote side of the connection.                                                                                                                                                                          |
| <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

```

## show interfaces (Aggregated Ethernet)

### Syntax

```
show interfaces aenumber
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Release Information

Command introduced before Junos OS Release 7.4.

### Description

(M Series, T Series, MX Series, and PTX Series routers) Display status information about the specified aggregated Ethernet interfaces.

### Options

***aenumber***—Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.

**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.

**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

### RELATED DOCUMENTATION

**List of Sample Output**

[show interfaces \(Aggregated Ethernet\) on page 1561](#)

[show interfaces brief \(Aggregated Ethernet\) on page 1562](#)

[show interfaces detail \(Aggregated Ethernet\) on page 1562](#)

[show interfaces extensive \(Aggregated Ethernet\) on page 1563](#)

[show interfaces extensive \(Aggregated Ethernet with VLAN Stacking\) on page 1565](#)

**Output Fields**

Table 135 on page 1554 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

**Table 135: 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                      |
| <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                      |

Table 135: Aggregated Ethernet show interfaces Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Output         |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <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              |
| <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: 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> |
| <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 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 135: 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 aborts 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> |
| <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 aborts 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> |

Table 135: Aggregated Ethernet show interfaces 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 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>         |
| <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 |
| <b>SNMP ifIndex</b>            | SNMP interface index number of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b><br>none |

Table 135: Aggregated Ethernet show interfaces 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 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 135: 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> <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> </li> </ul> | <b>detail extensive</b><br>none |



Table 135: 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> |                                 |
| 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>NOTE: 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> | <p>Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>brief</b>                    |
| <b>Protocol</b>        | <p>Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive</b><br>none |
| <b>MTU</b>             | <p>Maximum transmission unit size on the logical interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b><br>none |

Table 135: Aggregated Ethernet show interfaces Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                          | Level of Output                 |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <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>         |
| <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: 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   :                60                0 bps
    Output bytes   :                 0                0 bps
    Input  packets :                 1                0 pps
    Output packets :                 0                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: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
  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  :      1        0        60        0
    Output:      0        0         0        0
  Adaptive Statistics:
    Adaptive Adjusts:      0
    Adaptive Scans  :      0
    Adaptive Updates:      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 :           1           0           60           0
      Output:           0           0           0           0
LACP info:      Role      System              System      Port      Port      Port
                priority      identifier  priority  number      key

fe-1/0/3.0      Actor      127  00:00:5e:00:53:85      127      2      1
fe-1/0/3.0      Partner    127  00:00:5e:00:53:c3      127      1      1
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
fe-1/0/3.0              3188          3186              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: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
  Generation: 49

```

### 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 demux0 (Demux Interfaces)

### Syntax

```
show interfaces demux0.logical-interface-number
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Release Information

Command introduced in Junos OS Release 9.0.

### Description

(MX Series and M Series routers only) Display status information about the specified demux interface.

### Options

**none**—Display standard information about the specified demux interface.

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

**descriptions**—(Optional) Display interface description strings.

**media**—(Optional) Display media-specific information about network interfaces.

**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

### RELATED DOCUMENTATION

*Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

### List of Sample Output

[show interfaces demux0 \(Demux\) on page 1576](#)

[show interfaces demux0 \(PPPoE over Aggregated Ethernet\) on page 1577](#)

[show interfaces demux0 extensive \(Targeted Distribution for Aggregated Ethernet Links\) on page 1578](#)

[show interfaces demux0 \(ACI Interface Set Configured\) on page 1578](#)

### Output Fields

Table 136 on page 1569 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 136: 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</b> none |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                  | <b>brief detail</b><br><b>extensive</b> none |
| <b>Enabled</b>            | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | <b>brief detail</b><br><b>extensive</b> none |
| <b>Physical link</b>      | Status of the physical link ( <b>Up</b> or <b>Down</b> ).                                                                            | <b>detail</b> <b>extensive</b><br>none       |
| <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</b> <b>extensive</b><br>none       |
| <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>                                 |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                        | <b>detail</b> <b>extensive</b><br>none       |
| <b>Generation</b>         | Unique number for use by Juniper Networks technical support only.                                                                    | <b>detail</b> <b>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>      |

Table 136: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                    | Level of Output                              |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| <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 <i>Common Output Fields Description</i> .                                                                                                   | <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 <i>Common Output Fields Description</i> .                                                                                                      | <b>brief detail</b><br><b>extensive</b> none |
| <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><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: year-month-day hour:minute:second:timezone (hour:minute:second 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>                      |

Table 136: 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> |
| <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 aborts 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 136: 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 aborts 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              |
| <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 <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>brief detail</b><br><b>extensive</b> none |
| <b>Encapsulation</b>                 | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>brief extensive</b> none                  |
| <b>ACI VLAN:<br/>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 |

Table 136: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output                 |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Demux</b>                  | <p>Specific IP demultiplexing (demux) values:</p> <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—in</b>et family prefix.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b><br>none |
| <b><i>protocol-family</i></b> | Protocol family configured on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>brief</b>                    |
| <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 136: 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 <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>             | Route table in which the logical interface address is located. For example, 0 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 <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b><br>none |

Table 136: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                                                                                                    | Level of Output                                   |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| <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                   |
| <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>                                      |
| <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                   |



Table 136: show interfaces demux0 (Demux Interfaces) Output Fields (*continued*)

| Field Name | Field Description                | Level of Output          |
|------------|----------------------------------|--------------------------|
| AC Name    | Name of the access concentrator. | detail extensive<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

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

### Syntax

```
show interfaces diagnostics optics interface-name
```

### 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.

### 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-fpc/pic/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-fpc/pic/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 207.

#### Required Privilege Level

view

#### RELATED DOCUMENTATION

*Determining Transceiver Support and Specifications*

#### List of Sample Output

[show interfaces diagnostics optics \(DWDM and DWDM OTN\) on page 1604](#)

[show interfaces diagnostics optics \(MPC6E with OTN MIC\) on page 1605](#)

[show interfaces diagnostics optics \(Bidirectional SFP\) on page 1606](#)

[show interfaces diagnostics optics \(SFP\) on page 1607](#)

[show interfaces diagnostics optics \(SFP\) on page 1608](#)

[show interfaces diagnostics optics \(XFP and CFP Optics\) on page 1609](#)

[show interfaces diagnostics optics for 10-Gigabit Ethernet \(PTX 24-10GE-SFPP\) on page 1610](#)

[show interfaces diagnostics optics for 40-Gigabit Ethernet on page 1611](#)

[show interfaces diagnostics optics \(P1-PTX-2-100G-WDM\) on page 1615](#)

[show interfaces diagnostics optics \(P1-PTX-24-10G-W-SFPP \) on page 1617](#)

[show interfaces diagnostics optics \(P2-10G-40G-QSFPP PIC in 40-Gigabit Ethernet mode\) on page 1618](#)

[show interfaces diagnostics optics \(P2-10G-40G-QSFPP PIC in 10-Gigabit Ethernet mode\) on page 1620](#)

[show interfaces diagnostics optics \(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC\) on page 1621](#)

[show interfaces diagnostics optics \(PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC on page 1623](#)

[show interfaces diagnostics optics \(for VCP\) on page 1625](#)

[show interfaces diagnostics optics \(MPC7 with interfaces disabled\) on page 1626](#)

[show interfaces diagnostics optics for 100-Gigabit Ethernet \(QSFP-100GE-DWDM2\) on page 1629](#)

### Output Fields

[Table 137 on page 1582](#) 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 137: 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.                   |
| <b>Receiver signal average optical power</b> | 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. |
| <b>Laser end-of-life alarm</b>               | Laser end-of-life alarm: <b>On</b> or <b>Off</b> .                                                                                                                       |
| <b>Laser wavelength alarm</b>                | Laser wavelength alarm: <b>On</b> or <b>Off</b> .                                                                                                                        |
| <b>Laser bias current alarm</b>              | Laser bias current alarm: <b>On</b> or <b>Off</b> .                                                                                                                      |
| <b>Laser temperature alarm</b>               | Laser temperature alarm: <b>On</b> or <b>Off</b> .                                                                                                                       |
| <b>Laser power alarm</b>                     | Laser power alarm: <b>On</b> or <b>Off</b> .                                                                                                                             |
| <b>Modulator temperature alarm</b>           | Modulator temperature alarm: <b>On</b> or <b>Off</b> .<br><br>Transceivers from some vendors do not support this field.                                                  |
| <b>Modulator bias alarm</b>                  | Modulator bias alarm: <b>On</b> or <b>Off</b> .                                                                                                                          |
| <b>Tx multiplexer FIFO error alarm</b>       | Transmit multiplexer first in, first out (FIFO) error alarm: <b>On</b> or <b>Off</b> .                                                                                   |

Table 137: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet DWDM and DWDM OTN PICs (*continued*)

| Field Name                                    | Field Description                                                                                                          |
|-----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Tx loss of PLL lock alarm</b>              | Transmit loss of phase-locked loop (PLL) lock alarm: <b>On</b> or <b>Off</b> .                                             |
| <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><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 138 on page 1583 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 138: 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.                                                                                                                        |
| <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.                                                                                                                       |



Table 138: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics  
(continued)

| Field Name                             | Field Description                                                                |
|----------------------------------------|----------------------------------------------------------------------------------|
| <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> .                |
| <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> .              |

Table 138: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics  
(continued)

| Field Name                              | Field Description                                                                                |
|-----------------------------------------|--------------------------------------------------------------------------------------------------|
| Module temperature low warning          | Module temperature low warning. Displays <b>on</b> or <b>off</b> .                               |
| Module voltage high alarm               | Module voltage high alarm. Displays <b>on</b> or <b>off</b> .                                    |
| Module voltage low alarm                | Module voltage low alarm. Displays <b>on</b> or <b>off</b> .                                     |
| Module voltage high warning             | Module voltage high warning. Displays <b>on</b> or <b>off</b> .                                  |
| Module voltage low warning              | Module voltage high warning. Displays <b>on</b> or <b>off</b> .                                  |
| Laser rx power high alarm               | Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .                               |
| Laser rx power low alarm                | Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .                                |
| Laser rx power high warning             | Receive laser power high warning. Displays <b>on</b> or <b>off</b> .                             |
| Laser rx power low warning              | Receive laser power low warning. Displays <b>on</b> or <b>off</b> .                              |
| TEC fault alarm                         | TEC fault alarm. Displays <b>on</b> or <b>off</b> .                                              |
| Wavelength unlocked alarm               | Wavelength unlocked alarm. Displays <b>on</b> or <b>off</b> .                                    |
| TxTune                                  | Optical transmit side status. Displays whether optical transmit side is not ready due to tuning. |
| Laser bias current high alarm threshold | Vendor-specified threshold for the laser bias current high alarm: <b>70.000 mA</b> .             |
| Laser bias current low alarm threshold  | Vendor-specified threshold for the laser bias current low alarm: <b>0.0002 mA</b> .              |

Table 138: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics  
(continued)

| Field Name                                | Field Description                                                                                          |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------|
| Laser bias current high warning threshold | Vendor-specified threshold for the laser bias current high warning: <b>65.000 mA</b> .                     |
| Laser bias current low warning threshold  | Vendor-specified threshold for the laser bias current low warning: <b>0.0002 mA</b> .                      |
| Laser output power high alarm threshold   | Vendor-specified threshold for the laser output power high alarm: <b>1.0000 mW</b> or <b>0.00 dBm</b> .    |
| Laser output power low alarm threshold    | Vendor-specified threshold for the laser output power low alarm: <b>0.0560 mW</b> or <b>-12.52 dBm</b> .   |
| Laser output power high warning threshold | Vendor-specified threshold for the laser output power high warning: <b>0.6300 mW</b> or <b>-2.01 dBm</b> . |
| Laser output power low warning threshold  | Vendor-specified threshold for the laser output power low warning: <b>0.0890 mW</b> or <b>-10.51 dBm</b> . |
| Module temperature high alarm threshold   | Vendor-specified threshold for the module temperature high alarm: <b>100° C</b> or <b>212° F</b> .         |
| Module temperature low alarm threshold    | Vendor-specified threshold for the module temperature low alarm: <b>-50° C</b> or <b>-58° F</b> .          |
| Module temperature high warning threshold | Vendor-specified threshold for the module temperature high warning: <b>95 ° C</b> or <b>203 ° F</b> .      |
| Module temperature low warning threshold  | Vendor-specified threshold for the module temperature low warning: <b>-48° C</b> or <b>-54° F</b> .        |
| Module voltage high alarm threshold       | Module voltage high alarm threshold: <b>3.700 v</b> .                                                      |

Table 138: show interfaces diagnostics optics Output Fields for Gigabit Ethernet Bidirectional SFP Optics (continued)

| Field Name                            | Field Description                                                                       |
|---------------------------------------|-----------------------------------------------------------------------------------------|
| Module voltage low alarm threshold    | Module voltage low alarm threshold: 2.900 v.                                            |
| Module voltage high warning threshold | Module voltage high warning threshold: 3.7600 v.                                        |
| Module voltage low warning threshold  | Module voltage low warning threshold: 3.000 v.                                          |
| Laser rx power high alarm threshold   | Vendor-specified threshold for the laser Rx power high alarm: 1.9953 mW or 3.00 dBm.    |
| Laser rx power low alarm threshold    | Vendor-specified threshold for the laser Rx power low alarm: 0.0001 mW or -40.00 dBm.   |
| Laser rx power high warning threshold | Vendor-specified threshold for the laser Rx power high warning: 1.0000 mW or 0.00 dBm.  |
| Laser rx power low warning threshold  | Vendor-specified threshold for the laser Rx power low warning: 0.0010 mW or -30.00 dBm. |

Table 139 on page 1587 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 139: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers

| Field Name         | Field Description                       |
|--------------------|-----------------------------------------|
| Physical interface | Name of the physical interface.         |
| Laser bias current | Measured laser bias current in uA.      |
| Laser output power | Measured laser output power in mW.      |
| Module temperature | Internally measured module temperature. |
| Module voltage     | Internally measured module voltage.     |
| Laser rx power     | Measured receive optical power in mW.   |

Table 139: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (*continued*)

| Field Name                      | Field Description                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------|
| Laser bias current high alarm   | Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser bias current low alarm    | Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Laser output power high alarm   | Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser output power low alarm    | Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Module temp high alarm          | Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Module temp low alarm           | Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Laser rx power high alarm       | Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.    |
| Laser rx power low alarm        | Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser bias current high warning | Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| 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.  |

Table 139: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (*continued*)

| Field Name                                | Field Description                                                                               |
|-------------------------------------------|-------------------------------------------------------------------------------------------------|
| 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.             |
| Laser bias current high warning threshold | Laser bias current high warning threshold. Warning ranges are vendor-specific.                  |
| Laser bias current low warning threshold  | Laser bias current low warning threshold. Warning ranges are vendor-specific.                   |
| Laser output power high alarm threshold   | Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Laser output power low alarm threshold    | Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.             |
| Laser output power high warning threshold | Laser output power high warning threshold. Warning ranges are vendor-specific.                  |
| Laser output power low warning threshold  | Laser output power low warning threshold. Warning ranges are vendor-specific.                   |
| Module temperature high alarm threshold   | Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Module temperature low alarm threshold    | Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.             |

Table 139: show interfaces diagnostics Output Fields for Gigabit Ethernet SFP Transceivers (*continued*)

| Field Name                                | Field Description                                                                         |
|-------------------------------------------|-------------------------------------------------------------------------------------------|
| Module temperature high warning threshold | Module temperature high warning threshold. Warning ranges are vendor-specific.            |
| Module temperature low warning threshold  | Module temperature low warning threshold. Warning ranges are vendor-specific.             |
| Module voltage high alarm threshold       | Module voltage high alarm threshold. Alarm ranges are vendor-specific.                    |
| Module voltage low alarm threshold        | Module voltage low alarm threshold. Alarm ranges are vendor-specific.                     |
| Module voltage high warning threshold     | Module voltage high warning threshold. Warning ranges are vendor-specific.                |
| Module voltage low warning threshold      | Module voltage low warning threshold. Warning ranges are vendor-specific.                 |
| Laser rx power high alarm threshold       | Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.     |
| Laser rx power low alarm threshold        | Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.      |
| Laser rx power high warning threshold     | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific. |
| Laser rx power high low threshold         | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific. |

Table 140 on page 1590 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 140: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers

| Field Name         | Field Description                  |
|--------------------|------------------------------------|
| Physical interface | Name of the physical interface.    |
| Laser bias current | Measured laser bias current in mA. |

Table 140: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (continued)

| Field Name                      | Field Description                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------|
| Laser output power              | Measured laser output power in mW.                                                             |
| Module temperature              | Internally measured module temperature.                                                        |
| Laser rx power                  | Measured receive optical power in mW.                                                          |
| Laser bias current high alarm   | Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser bias current low alarm    | Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Laser output power high alarm   | Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser output power low alarm    | Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Module temp high alarm          | Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Module temp low alarm           | Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Laser rx power high alarm       | Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.    |
| Laser rx power low alarm        | Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser bias current high warning | Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| 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.  |



Table 140: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (*continued*)

| Field Name                                | Field Description                                                                               |
|-------------------------------------------|-------------------------------------------------------------------------------------------------|
| 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.             |
| Laser output power high alarm threshold   | Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Laser output power low alarm threshold    | Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.             |
| Module temperature high alarm threshold   | Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Module temperature low alarm threshold    | Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.             |
| Laser rx power high alarm threshold       | Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.           |
| Laser rx power low alarm threshold        | Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Laser bias current high warning threshold | Laser bias current high warning threshold. Warning ranges are vendor-specific.                  |

Table 140: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet Transceivers (*continued*)

| Field Name                                | Field Description                                                                         |
|-------------------------------------------|-------------------------------------------------------------------------------------------|
| Laser bias current low warning threshold  | Laser bias current low warning threshold. Warning ranges are vendor-specific.             |
| Laser output power high warning threshold | Laser output power high warning threshold. Warning ranges are vendor-specific.            |
| Laser output power low warning threshold  | Laser output power low warning threshold. Warning ranges are vendor-specific.             |
| Module temperature high warning threshold | Module temperature high warning threshold. Warning ranges are vendor-specific.            |
| Module temperature low warning threshold  | Module temperature low warning threshold. Warning ranges are vendor-specific.             |
| Laser rx power high warning threshold     | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific. |
| Laser rx power low warning threshold      | Laser receive power low warning threshold. Warning threshold ranges are vendor-specific.  |

Table 141 on page 1593 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 141: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers

| Field Name              | Field Description                                                                                                                                                                             |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical interface      | Name of the physical interface.                                                                                                                                                               |
| Grid Channel Number     | 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. |
| Corrected Error Ratio   | Indicates accumulated Bit Error Ratio.                                                                                                                                                        |
| Uncorrected Words Ratio | Monitors the error rate for either the full 100G link or on a channel-by-channel basis. Indicates the frame errors.                                                                           |

Table 141: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers (continued)

| Field Name                      | Field Description                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------|
| Laser bias current              | Measured laser bias current in mA.                                                             |
| Laser output power              | Measured laser output power in mW.                                                             |
| Module temperature              | Internally measured module temperature.                                                        |
| Laser rx power                  | Measured receive optical power in mW.                                                          |
| Laser bias current high alarm   | Laser bias current high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser bias current low alarm    | Laser bias current low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Laser output power high alarm   | Laser output power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser output power low alarm    | Laser output power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Module temp high alarm          | Module temperature high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Module temp low alarm           | Module temperature low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.      |
| Laser rx power high alarm       | Laser receive power high alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.    |
| Laser rx power low alarm        | Laser receive power low alarm: <b>On</b> or <b>Off</b> . Alarm ranges are vendor-specific.     |
| Laser bias current high warning | Laser bias current high warning: <b>On</b> or <b>Off</b> . Warning ranges are vendor-specific. |
| 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. |

Table 141: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers (continued)

| Field Name                              | Field Description                                                                               |
|-----------------------------------------|-------------------------------------------------------------------------------------------------|
| 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.             |
| Laser output power high alarm threshold | Laser output power high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Laser output power low alarm threshold  | Laser output power low alarm threshold. Alarm threshold ranges are vendor-specific.             |
| Module temperature high alarm threshold | Module temperature high alarm threshold. Alarm threshold ranges are vendor-specific.            |
| Module temperature low alarm threshold  | Module temperature low alarm threshold. Alarm threshold ranges are vendor-specific.             |
| Laser rx power high alarm threshold     | Laser receive power high alarm threshold. Alarm threshold ranges are vendor-specific.           |
| Laser rx power low alarm threshold      | Laser receive power low alarm threshold. Alarm threshold ranges are vendor-specific.            |

Table 141: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers (continued)

| Field Name                                | Field Description                                                                                                                                                                              |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Laser bias current high warning threshold | Laser bias current high warning threshold. Warning ranges are vendor-specific.                                                                                                                 |
| Laser bias current low warning threshold  | Laser bias current low warning threshold. Warning ranges are vendor-specific.                                                                                                                  |
| Laser output power high warning threshold | Laser output power high warning threshold. Warning ranges are vendor-specific.                                                                                                                 |
| Laser output power low warning threshold  | Laser output power low warning threshold. Warning ranges are vendor-specific.                                                                                                                  |
| Module temperature high warning threshold | Module temperature high warning threshold. Warning ranges are vendor-specific.                                                                                                                 |
| Module temperature low warning threshold  | Module temperature low warning threshold. Warning ranges are vendor-specific.                                                                                                                  |
| Laser rx power high warning threshold     | Laser receive power high warning threshold. Warning threshold ranges are vendor-specific.                                                                                                      |
| Laser rx power low warning threshold      | 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. |

Table 141: show interfaces diagnostics optics Output Fields for 100-Gigabit Ethernet Transceivers (continued)

| Field Name    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 142 on page 1597 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 142: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers

| Field Name                      | Field Description                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical interface              | Name of the physical interface.                                                                                                                            |
| Laser bias current              | Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| Laser output power              | 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.     |
| Module temperature              | Temperature of the XFP optics module, in Celsius and Fahrenheit.                                                                                           |
| Laser rx power                  | Laser received optical power, in mW and dBm.                                                                                                               |
| Laser bias current high alarm   | Laser bias power setting high alarm. Displays <b>on</b> or <b>off</b> .                                                                                    |
| Laser bias current low alarm    | Laser bias power setting low alarm. Displays <b>on</b> or <b>off</b> .                                                                                     |
| Laser bias current high warning | Laser bias power setting high warning. Displays <b>on</b> or <b>off</b> .                                                                                  |
| Laser bias current low warning  | Laser bias power setting low warning. Displays <b>on</b> or <b>off</b> .                                                                                   |
| Laser output power high alarm   | Laser output power high alarm. Displays <b>on</b> or <b>off</b> .                                                                                          |

Table 142: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers  
(continued)

| Field Name                      | Field Description                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Laser output power low alarm    | Laser output power low alarm. Displays <b>on</b> or <b>off</b> .                                                                          |
| Laser output power high warning | Laser output power high warning. Displays <b>on</b> or <b>off</b> .                                                                       |
| Laser output power low warning  | Laser output power low warning. Displays <b>on</b> or <b>off</b> .                                                                        |
| Module temperature high alarm   | Module temperature high alarm. Displays <b>on</b> or <b>off</b> .                                                                         |
| Module temperature low alarm    | Module temperature low alarm. Displays <b>on</b> or <b>off</b> .                                                                          |
| Module temperature high warning | Module temperature high warning. Displays <b>on</b> or <b>off</b> .                                                                       |
| Module temperature low warning  | Module temperature low warning. Displays <b>on</b> or <b>off</b> .                                                                        |
| Laser rx power high alarm       | Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .                                                                        |
| Laser rx power low alarm        | Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .                                                                         |
| Laser rx power high warning     | Receive laser power high warning. Displays <b>on</b> or <b>off</b> .                                                                      |
| Laser rx power low warning      | Receive laser power low warning. Displays <b>on</b> or <b>off</b> .                                                                       |
| Module not ready alarm          | Module not ready alarm. When <b>on</b> , indicates the module has an operational fault. Displays <b>on</b> or <b>off</b> .                |
| Module power down alarm         | 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> . |

Table 142: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers  
(continued)

| Field Name                                       | Field Description                                                                                                                           |
|--------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <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> . |
| <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> .                                     |



**Table 142: show interfaces diagnostics optics Output Fields for 10-Gigabit Ethernet XFP Transceivers**  
(continued)

| Field Name                                | Field Description                                                                                         |
|-------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| Laser output power high warning threshold | Vendor-specified threshold for the laser output power high warning: <b>0.7940 mW</b> or <b>-100 dBm</b> . |
| Laser output power low warning threshold  | Vendor-specified threshold for the laser output power low warning: <b>0.2510 mW</b> or <b>-600 dBm</b> .  |
| Module temperature high alarm threshold   | Vendor-specified threshold for the module temperature high alarm: <b>90° C</b> or <b>194° F</b> .         |
| Module temperature low alarm threshold    | Vendor-specified threshold for the module temperature low alarm: <b>-5° C</b> or <b>23° F</b> .           |
| Module temperature high warning threshold | Vendor-specified threshold for the module temperature high warning: <b>85° C</b> or <b>185° F</b> .       |
| Module temperature low warning threshold  | Vendor-specified threshold for the module temperature low warning: <b>0° C</b> or <b>32° F</b> .          |
| Laser rx power high alarm threshold       | Vendor-specified threshold for the laser Rx power high alarm: <b>1.2589 mW</b> or <b>1.00 dBm</b> .       |
| Laser rx power low alarm threshold        | Vendor-specified threshold for the laser Rx power low alarm: <b>0.0323 mW</b> or <b>-14.91 dBm</b> .      |
| Laser rx power high warning threshold     | Vendor-specified threshold for the laser Rx power high warning: <b>1.1220 mW</b> or <b>0.50 dBm</b> .     |
| Laser rx power low warning threshold      | Vendor-specified threshold for the laser Rx power low warning: <b>0.0363 mW</b> or <b>-14.40 dBm</b> .    |

[Table 143 on page 1601](#) 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 143: show interfaces diagnostics optics Output for Virtual Chassis Ports

| Field Name                            | Field Description                                                                                                                                          |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical interface                    | Name of the physical interface.                                                                                                                            |
| Laser bias current                    | Magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents. |
| Laser output power                    | Laser output power, in milliwatts (mW) and decibels, referenced to 1.0 mW (dBm).                                                                           |
| Module temperature                    | Temperature of the optics module, in Celsius and Fahrenheit.                                                                                               |
| Module voltage                        | Internally measured module voltage.                                                                                                                        |
| Receiver signal average optical power | Average received optical power, in mW and dBm.                                                                                                             |
| Laser bias current high alarm         | Laser bias power setting high alarm. Displays <b>on</b> or <b>off</b> .                                                                                    |
| Laser bias current low alarm          | Laser bias power setting low alarm. Displays <b>on</b> or <b>off</b> .                                                                                     |
| Laser bias current high warning       | Laser bias power setting high warning. Displays <b>on</b> or <b>off</b> .                                                                                  |
| Laser bias current low warning        | Laser bias power setting low warning. Displays <b>on</b> or <b>off</b> .                                                                                   |
| Laser output power high alarm         | Laser output power high alarm. Displays <b>on</b> or <b>off</b> .                                                                                          |
| Laser output power low alarm          | Laser output power low alarm. Displays <b>on</b> or <b>off</b> .                                                                                           |
| Laser output power high warning       | Laser output power high warning. Displays <b>on</b> or <b>off</b> .                                                                                        |
| Laser output power low warning        | Laser output power low warning. Displays <b>on</b> or <b>off</b> .                                                                                         |
| Module temperature high alarm         | Module temperature high alarm. Displays <b>on</b> or <b>off</b> .                                                                                          |

Table 143: show interfaces diagnostics optics Output for Virtual Chassis Ports (continued)

| Field Name                              | Field Description                                                    |
|-----------------------------------------|----------------------------------------------------------------------|
| Module temperature low alarm            | Module temperature low alarm. Displays <b>on</b> or <b>off</b> .     |
| Module temperature high warning         | Module temperature high warning. Displays <b>on</b> or <b>off</b> .  |
| Module temperature low warning          | Module temperature low warning. Displays <b>on</b> or <b>off</b> .   |
| Module voltage high alarm               | Module voltage high alarm. Displays <b>on</b> or <b>off</b> .        |
| Module voltage low alarm                | Module voltage low alarm. Displays <b>on</b> or <b>off</b> .         |
| Module voltage high warning             | Module voltage high warning. Displays <b>on</b> or <b>off</b> .      |
| Module voltage low warning              | Module voltage high warning. Displays <b>on</b> or <b>off</b> .      |
| Laser rx power high alarm               | Receive laser power high alarm. Displays <b>on</b> or <b>off</b> .   |
| Laser rx power low alarm                | Receive laser power low alarm. Displays <b>on</b> or <b>off</b> .    |
| Laser rx power high warning             | Receive laser power high warning. Displays <b>on</b> or <b>off</b> . |
| Laser rx power low warning              | Receive laser power low warning. Displays <b>on</b> or <b>off</b> .  |
| Laser bias current high alarm threshold | Vendor-specified threshold for the laser bias current high alarm.    |
| Laser bias current low alarm threshold  | Vendor-specified threshold for the laser bias current low alarm.     |

Table 143: show interfaces diagnostics optics Output for Virtual Chassis Ports *(continued)*

| Field Name                                | Field Description                                                   |
|-------------------------------------------|---------------------------------------------------------------------|
| Laser bias current high warning threshold | Vendor-specified threshold for the laser bias current high warning. |
| Laser bias current low warning threshold  | Vendor-specified threshold for the laser bias current low warning.  |
| Laser output power high alarm threshold   | Vendor-specified threshold for the laser output power high alarm.   |
| Laser output power low alarm threshold    | Vendor-specified threshold for the laser output power low alarm.    |
| Laser output power high warning threshold | Vendor-specified threshold for the laser output power high warning. |
| Laser output power low warning threshold  | Vendor-specified threshold for the laser output power low warning.  |
| Module temperature high alarm threshold   | Vendor-specified threshold for the module temperature high alarm.   |
| Module temperature low alarm threshold    | Vendor-specified threshold for the module temperature low alarm.    |
| Module temperature high warning threshold | Vendor-specified threshold for the module temperature high warning. |
| Module temperature low warning threshold  | Vendor-specified threshold for the module temperature low warning.  |
| Module voltage high alarm threshold       | Module voltage high alarm threshold.                                |
| Module voltage low alarm threshold        | Module voltage low alarm threshold.                                 |

Table 143: show interfaces diagnostics optics Output for Virtual Chassis Ports (continued)

| Field Name                            | Field Description                                               |
|---------------------------------------|-----------------------------------------------------------------|
| Module voltage high warning threshold | Module voltage high warning threshold.                          |
| Module voltage low warning threshold  | Module voltage low warning threshold.                           |
| Laser rx power high alarm threshold   | Vendor-specified threshold for the laser Rx power high alarm.   |
| Laser rx power low alarm threshold    | Vendor-specified threshold for the laser Rx power low alarm.    |
| Laser rx power high warning threshold | Vendor-specified threshold for the laser Rx power high warning. |
| Laser rx power low warning threshold  | Vendor-specified threshold for the laser Rx power low warning.  |

## Sample Output

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 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

```

# show interfaces (far-end-interval)

## Syntax

```
show interfaces far-end-interval interface-fpc/pic/port
```

## Release Information

Command introduced in Junos OS Release 9.4.

## Description

On channelized interfaces, display the far end interval data for the specified interface.

## Required Privilege Level

view

## List of Sample Output

[show interfaces far-end-interval coc12-5/2/0 on page 1631](#)

[show interfaces far-end-interval coc1-5/2/1:1 on page 1632](#)

## Output Fields

[Table 144 on page 1631](#) 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 144: show interfaces far-end-interval Output Fields

| Field Name         | Field Description                         |
|--------------------|-------------------------------------------|
| Physical interface | Interface FPC/PIC/port values.            |
| SNMP ifIndex       | SNMP interface index value.               |
| ES-L/P             | Error detection—Errored seconds.          |
| SES-L/P            | Error detection—Severely errored seconds. |
| UAS-L/P            | 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:

```

## show interfaces (Fast Ethernet)

### Syntax

```
show interfaces interface-type
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Release Information

Command introduced before Junos OS Release 7.4.

### Description

Display status information about the specified Fast Ethernet interface.

### Options

***interface-type***—On M Series and T Series routers, the interface type is ***fe-fpc/pic/port***.

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

**descriptions**—(Optional) Display interface description strings.

**media**—(Optional) Display media-specific information about network interfaces.

**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

### List of Sample Output

[show interfaces \(Fast Ethernet\) on page 1651](#)

[show interfaces brief \(Fast Ethernet\) on page 1652](#)

[show interfaces detail \(Fast Ethernet\) on page 1652](#)

[show interfaces extensive \(Fast Ethernet\) on page 1653](#)

### Output Fields

[Table 145 on page 1634](#) lists the output fields for the **show interfaces** (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 145: 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 <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 none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                                                                                             | <b>detail extensive none</b> |
| <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: <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                   |
| <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                   |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                    | Level of Output                 |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <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.                                                                                                                                                                                                                     | <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: year-month-day hour:minute:second:timezone (hour:minute:second 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>Input Rate</b>       | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                    | None specified                  |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output         |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <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 145: 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 aborts 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 145: 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 aborts 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> <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> |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <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> |
| <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>        |



Table 145: show interfaces Fast 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 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 extensive</b><br>none |
| <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>—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>                                                                                                                                                                                                                                                                                                                                               |                                 |
| <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 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| MAC statistics | <p>Receive and Transmit 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> <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> | extensive       |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                            | Field Description                                                                    | Level of Output  |
|---------------------------------------|--------------------------------------------------------------------------------------|------------------|
| <b>OTN Received Overhead Bytes</b>    | APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58<br>Payload Type: 0x08 | <b>extensive</b> |
| <b>OTN Transmitted Overhead Bytes</b> | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00<br>Payload Type: 0x08 | <b>extensive</b> |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output  |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Filter statistics</b> | <p><b>Receive and 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 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.</li> <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>extensive</b> |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <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> |
| <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 145: 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 145: 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 145: 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> <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>extensive</b> |



Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                 |
|--------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <b>Received path trace,<br/>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 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <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>                |
| <b>CoS information</b>                                 | Information about the CoS queue for the physical interface. <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 145: 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 <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 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 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 extensive</b> none          |
| <b>MTU</b>           | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive</b> none          |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <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>                |
| <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> 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> 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> none    |

Table 145: show interfaces Fast Ethernet Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                    | Level of Output                 |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <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>         |
| <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 <i>Common Output Fields Description</i> .                                                  | <b>detail extensive</b> 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 extensive</b> none    |
| <b>Destination</b>            | IP address of the remote side of the connection.                                                                                                                                                     | <b>detail extensive</b> none    |
| <b>Local</b>                  | IP address of the logical interface.                                                                                                                                                                 | <b>detail extensive</b> none    |
| <b>Broadcast</b>              | Broadcast address of the logical interlace.                                                                                                                                                          | <b>detail extensive</b> 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          0
Jabber frames   0          0
Fragment frames 0          0
VLAN tagged frames 0        0
Code violations 0          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



## show interfaces

### List of Syntax

[Syntax \(Gigabit Ethernet\) on page 1656](#)

[Syntax \(10 Gigabit Ethernet\) on page 1656](#)

[Syntax \(SRX Series Devices and \(vSRX and vSRX 3.0 platforms\)\) on page 1656](#)

### 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 (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 l2-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>
)

```

### Release Information

Command introduced before Junos OS Release 7.4 for Gigabit interfaces.

Command introduced in Junos OS Release 8.0 for 10 Gigabit interfaces.

Command modified in Junos OS Release 9.5 for SRX Series devices.

Command introduced in Junos OS Release 18.1 for Gigabit interfaces.

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.

### 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.

**NOTE:** 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.

### 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.

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

**descriptions**—(Optional) Display interface description strings.

**media**—(Optional) Display media-specific information about network interfaces.

**snmp-index *snmp-index***—(Optional) Display information for the specified SNMP index of the interface.

**statistics**—(Optional) Display static interface statistics.

For 10 Gigabit interfaces:

**xe-*fpc/pic/port***—Display standard information about the specified 10-Gigabit Ethernet interface.

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

**descriptions**—(Optional) Display interface description strings.

**media**—(Optional) Display media-specific information about network interfaces.

**snmp-index *snmp-index***—(Optional) Display information for the specified SNMP index of the interface.

**statistics**—(Optional) Display static interface statistics.

For SRX interfaces:

- **interface-name**—(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.

- **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
- **ce1-*pim*/0/ *port***—Channelized E1 interface.
- **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
- **ct1-*pim*/0/*port***—Channelized T1 interface.
- **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).
- **interface-name**—(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.
  - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
  - **ce1-*pim*/0/ *port***—Channelized E1 interface.
  - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
  - **ct1-*pim*/0/*port***—Channelized T1 interface.
  - **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

**Release History Table**

| Release                | Description                                                                                                                                                                                                      |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">18.4R1</a> | Starting in Junos OS Release 18.4R1, Output fields <b>Next-hop</b> and <b>vpls-status</b> is displayed in the <b>show interfaces interface name detail</b> 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](#)
**List of Sample Output**
[show interfaces \(Gigabit Ethernet\) on page 1709](#)
[show interfaces \(Gigabit Ethernet on MX Series Routers\) on page 1710](#)
[show interfaces \(link degrade status\) on page 1711](#)
[show interfaces extensive \(Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration\) on page 1712](#)
[show interfaces brief \(Gigabit Ethernet\) on page 1713](#)
[show interfaces detail \(Gigabit Ethernet\) on page 1713](#)
[show interfaces extensive \(Gigabit Ethernet IQ2\) on page 1715](#)
[show interfaces \(Gigabit Ethernet Unnumbered Interface\) on page 1719](#)
[show interfaces \(ACI Interface Set Configured\) on page 1719](#)
[show interfaces \(ALI Interface Set\) on page 1720](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, IQ2\) on page 1720](#)
[show interfaces extensive \(10-Gigabit Ethernet, WAN PHY Mode\) on page 1723](#)
[show interfaces extensive \(10-Gigabit Ethernet, DWDM OTN PIC\) on page 1726](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode\) on page 1729](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only\) on page 1730](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only\) on page 1731](#)
[Sample Output SRX Gigabit Ethernet on page 1733](#)
[Sample Output SRX Gigabit Ethernet on page 1733](#)
[show interfaces \(Gigabit Ethernet for vSRX and vSRX 3.0\) on page 1734](#)
[show interfaces detail \(Gigabit Ethernet\) on page 1735](#)
[show interfaces statistics st0.0 detail on page 1737](#)

[show interfaces extensive \(Gigabit Ethernet\) on page 1738](#)  
[show interfaces terse on page 1741](#)  
[show interfaces terse \(vSRX and vSRX 3.0\) on page 1742](#)  
[show interfaces controller \(Channelized E1 IQ with Logical E1\) on page 1743](#)  
[show interfaces controller \(Channelized E1 IQ with Logical DS0\) on page 1743](#)  
[show interfaces descriptions on page 1743](#)  
[show interfaces destination-class all on page 1744](#)  
[show interfaces diagnostics optics on page 1744](#)  
[show interfaces far-end-interval coc12-5/2/0 on page 1745](#)  
[show interfaces far-end-interval coc1-5/2/1:1 on page 1746](#)  
[show interfaces filters on page 1746](#)  
[show interfaces flow-statistics \(Gigabit Ethernet\) on page 1747](#)  
[show interfaces interval \(Channelized OC12\) on page 1748](#)  
[show interfaces interval \(E3\) on page 1749](#)  
[show interfaces interval \(SONET/SDH\) \(SRX devices\) on page 1749](#)  
[show interfaces load-balancing \(SRX devices\) on page 1750](#)  
[show interfaces load-balancing detail \(SRX devices\) on page 1750](#)  
[show interfaces mac-database \(All MAC Addresses on a Port SRX devices\) on page 1750](#)  
[show interfaces mac-database \(All MAC Addresses on a Service SRX devices\) on page 1751](#)  
[show interfaces mac-database mac-address on page 1752](#)  
[show interfaces mc-ae \(SRX devices\) on page 1752](#)  
[show interfaces media \(SONET/SDH\) on page 1752](#)  
[show interfaces policers \(SRX devices\) on page 1753](#)  
[show interfaces policers interface-name \(SRX devices\) on page 1754](#)  
[show interfaces queue \(SRX devices\) on page 1754](#)  
[show interfaces redundancy \(SRX devices\) on page 1755](#)  
[show interfaces redundancy \(Aggregated Ethernet SRX devices\) on page 1755](#)  
[show interfaces redundancy detail \(SRX devices\) on page 1756](#)  
[show interfaces routing brief \(SRX devices\) on page 1756](#)  
[show interfaces routing detail \(SRX devices\) on page 1757](#)  
[show interfaces routing-instance all \(SRX devices\) on page 1757](#)  
[show interfaces snmp-index \(SRX devices\) on page 1758](#)  
[show interfaces source-class all \(SRX devices\) on page 1758](#)  
[show interfaces statistics \(Fast Ethernet SRX devices\) on page 1759](#)  
[show interfaces switch-port \(SRX devices\) on page 1760](#)  
[show interfaces transport pm \(SRX devices\) on page 1760](#)  
[show security zones \(SRX devices\) on page 1762](#)

## Output Fields

[Table 146 on page 1662](#) 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 147 on page 1700](#).

Table 146: 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                   |
| <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 none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                                                                                             | <b>detail extensive none</b> |
| <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                   |
| <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                   |

Table 146: show interfaces (Gigabit Ethernet) 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. 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.                                                                                                                                                                                                                                            | <b>detail extensive</b> 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> none |
| <b>Hardware address</b> | Hardware MAC address.                                                                                                                                                                                                                                                       | <b>detail extensive</b> none |
| <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 hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                        | <b>detail extensive</b> 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                         |



Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Level of Output         |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <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> |
| <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</a> command.</p> | <b>detail extensive</b> |

Table 146: 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 aborts 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 146: show interfaces (Gigabit Ethernet) 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>• <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 aborts 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> <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> | extensive       |

Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Level of Output         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <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> <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>           | <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>        |

Table 146: 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 none</b> |
| 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>      |
| <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>Errorred 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 146: 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 146: show interfaces (Gigabit Ethernet) Output Fields (continued)

| Field Name     | Field Description | Level of Output |
|----------------|-------------------|-----------------|
| MAC statistics |                   | extensive       |

Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
|            | <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>—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 (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> <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> |                 |



Table 146: show interfaces (Gigabit Ethernet) Output Fields (continued)

| Field Name                            | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output  |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
|                                       | <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<br>Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b> |
| <b>OTN Transmitted Overhead Bytes</b> | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00<br>Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>extensive</b> |

Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Level of Output  |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <b>Filter statistics</b> | <p><b>Receive and 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 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.</li> <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>extensive</b> |

Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Level of Output  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <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> |
| <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 146: 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 146: 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 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Autonegotiation information | <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>—OK 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> <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> <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> | extensive       |

Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name                                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output              |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <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>             |
| <b>CoS information</b>                             | Information about the CoS queue for the physical interface. <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 none</b> |
| <b>SNMP ifIndex</b>                                | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |
| <b>Generation</b>                                  | Unique number for use by Juniper Networks technical support only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>detail extensive</b>      |

Table 146: 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 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 extensive</b> none          |
| <b>Encapsulation</b> | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels                            |



Table 146: show interfaces (Gigabit Ethernet) Output Fields (continued)

| Field Name                                         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output             |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------|
| ACI VLAN                                           | <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>                                                                                                                                                                                                                   | brief detail extensive none |
| 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> | detail                      |
| Protocol                                           | Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | detail extensive none       |
| MTU                                                | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail extensive none       |
| Neighbor Discovery Protocol (NDP) Queue Statistics | <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                  |

Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output              |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <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 none</b> |
| <b>Service Name Table</b>      | Name of the service name table for the interface configured with a PPPoE family.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive none</b> |
| <b>Max Sessions</b>            | Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail extensive none</b> |
| <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 none</b> |
| Direct Connect                 | 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 none</b> |
| <b>AC Name</b>                 | Name of the access concentrator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive none</b> |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>detail extensive none</b> |
| <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 146: 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 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 extensive</b>      |
| <b>Donor interface</b>          | (Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive 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 extensive 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 parentheses next to all interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail 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 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 none</b> |
| <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 none</b> |

Table 146: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                  | Level of Output              |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <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 the address flag. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | <b>detail extensive none</b> |
| <b>Destination</b>            | IP address of the remote side of the connection.                                                                                                   | <b>detail extensive none</b> |
| <b>Local</b>                  | IP address of the logical interface.                                                                                                               | <b>detail extensive none</b> |
| <b>Broadcast</b>              | Broadcast address of the logical interface.                                                                                                        | <b>detail extensive none</b> |
| <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 extensive none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                        | <b>detail extensive none</b> |
| <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                             |
| <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.                                                                                                                                                                                                                    | <b>detail</b><br><b>extensive</b> 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.                                                                                                                                                                                 | <b>detail</b><br><b>extensive</b>      |

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                          |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| Current address          | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | detail<br>extensive none |
| Hardware address         | Hardware MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | detail<br>extensive none |
| Last flapped             | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | detail<br>extensive none |
| Input Rate               | 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 specified           |
| Output Rate              | 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           |
| Statistics last cleared  | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | detail<br>extensive      |
| Egress account overhead  | Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | detail<br>extensive      |
| Ingress account overhead | Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | detail<br>extensive      |
| Traffic statistics       | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <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> | detail<br>extensive      |

**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 aborts and FCS errors.
- **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.
- **Framing errors**—Number of packets received with an invalid frame checksum (FCS).
- **Runts**—Number of frames received that are smaller than the runt threshold.
- **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.
- **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (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:

**extensive**

- **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.
- **Errors**—Sum of the outgoing frame aborts and FCS errors.
- **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.
- **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.
- **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.
- **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.
- **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
- **MTU errors**—Number of packets whose size exceeded the MTU of the interface.
- **Resource errors**—Sum of transmit drops.



|                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                  |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| <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<br/>extensive none</b> |
| <b>OTN alarms</b>                       | <p>Active OTN alarms identified on the interface.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>detail<br/>extensive</b>      |

|                    |                                                                                                                                                                                                                                                                                                                                                                               |                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| OTN defects        | OTN defects received on the interface.                                                                                                                                                                                                                                                                                                                                        | detail<br>extensive |
| OTN FEC Mode       | <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> | detail<br>extensive |
| 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  | Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                                                        | 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> |
| <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<br/>extensive</b> |

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |           |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| MAC statistics                 | <p>Receive and Transmit 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.</li> <li>• <b>Unicast packets, Broadcast packets, and 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> <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 runs (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> | extensive |
| OTN Received Overhead Bytes    | APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | extensive |
| OTN Transmitted Overhead Bytes | APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | extensive |

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |           |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Filter statistics | <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 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.</li> <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> | extensive |
| PMA PHY           | <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>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | extensive |

|             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |           |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| WIS section | <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>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | extensive |
| WIS line    | <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> | extensive |

|          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |           |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| WIS path | <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 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> | extensive |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|

|                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |           |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Autonegotiation information                    | <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> <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> | extensive |
| Received path trace,<br>Transmitted path trace | <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>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | extensive |



|                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |                  |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <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> |
| <b>CoS information</b>                        | Information about the CoS queue for the physical interface. <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> none |
| <b>SNMP ifIndex</b>      | SNMP interface index number for the logical interface.                                                                                                          | <b>detail</b><br><b>extensive</b> 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> 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> 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> none       |
| <b>MTU</b>            | Maximum transmission unit size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>detail</b><br><b>extensive</b> 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> 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<br/>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>       | <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<br/>extensive</b>      |
| <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<br/>extensive 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 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 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> 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> 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> none |
| <b>Destination</b>            | IP address of the remote side of the connection.                                                                                                                                                     | <b>detail</b><br><b>extensive</b> none |
| <b>Local</b>                  | IP address of the logical interface.                                                                                                                                                                 | <b>detail</b><br><b>extensive</b> none |
| <b>Broadcast</b>              | Broadcast address of the logical interlace.                                                                                                                                                          | <b>detail</b><br><b>extensive</b> 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 147: 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                                                                                                                 |                                                                                                                                                    |
| Outbound physical interface | <b>show interfaces ge-0/0/0 extensive</b>    | Traffic statistics:<br><br>Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes<br><br>MAC statistics:<br><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><br>Input bytes: 478 bytes per packet, representing the Layer 3 packet                                                                                                                 |                                                                                                                                                    |

[Table 148 on page 1701](#) lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 148: 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                   |
| <b>Interface index</b>    | Index number of the physical interface, which reflects its initialization sequence.                                                                                                                                                                 | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                                                                                                                       | <b>detail extensive none</b> |
| <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                   |
| <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                   |

Table 148: show interfaces Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output       |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| CoS queues                       | Number of CoS queues configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | detail extensive none |
| Current address                  | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | detail extensive none |
| Last flapped                     | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                          | detail extensive none |
| Input Rate                       | Input rate in bits per second (bps) and packets per second (pps).                                                                                                                                                                                                                                                                                                                                                                                                                                                             | None                  |
| Output Rate                      | Output rate in bps and pps.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | None                  |
| Active alarms and Active defects | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. 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> | detail extensive none |
| Statistics last cleared          | Time when the statistics for the interface were last set to zero.                                                                                                                                                                                                                                                                                                                                                                                                                                                             | detail extensive      |
| Traffic statistics               | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <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>                                                                  | detail extensive      |

Table 148: 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 the incoming frame aborts 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>.</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 148: 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 aborts 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> <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>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 148: show interfaces Output Fields (*continued*)

| Field Name     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Level of Output |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| MAC statistics | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> <li>• <b>Total octets and total packets</b>—Total number of octets and packets.</li> <li>• <b>Unicast packets, Broadcast packets, and 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> <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> | extensive       |

Table 148: 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 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.</li> <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>extensive</b> |

Table 148: show interfaces 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> </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>        |
| <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> |

Table 148: show interfaces Output Fields (*continued*)

| 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 none</b> |
| <b>SNMP ifIndex</b>          | SNMP interface index number for the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>detail extensive none</b> |
| <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>             |
| <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>             |

Table 148: show interfaces Output Fields (*continued*)

| Field Name                                            | Field Description                                                                                                 | Level of Output              |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|------------------------------|
| <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 none</b> |
| <b>MTU</b>                                            | Maximum transmission unit size on the logical interface.                                                          | <b>detail extensive none</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, 0 refers to the routing table inet.0. | <b>detail extensive none</b> |
| <b>Flags</b>                                          | Information about protocol family flags. .                                                                        | <b>detail extensive</b>      |
| <b>Addresses, Flags</b>                               | Information about the address flags..                                                                             | <b>detail extensive none</b> |
| <b>Destination</b>                                    | IP address of the remote side of the connection.                                                                  | <b>detail extensive none</b> |
| <b>Local</b>                                          | IP address of the logical interface.                                                                              | <b>detail extensive none</b> |
| <b>Broadcast</b>                                      | Broadcast address of the logical interface.                                                                       | <b>detail extensive none</b> |
| <b>Generation</b>                                     | Unique number for use by Juniper Networks technical support only.                                                 | <b>detail extensive</b>      |

## Sample Output Gigabit Ethernet

### 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: Sendbroadcast-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
        Bit errors          Seconds
        Bit errors          0
        Errored blocks      0
    Link Degraded* :
    Link Monitoring  : Enable

```



```

Link Degrade Set Threshold:      : 1E-7
Link Degrade Clear Threshold:    : 1E-12
Estimated BER                   : 1E-7
Link-degrade event              : Seconds      Count      State
                                782            1        Defect Active

```

**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:    3538446506             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              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 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 146 on page 1662](#).

### 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
    Policers: 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 00 00 00 00 00 00 00 00 00 00 00 00 00 00 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:
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
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 00 00 .....
Received SAPI:
00 00 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 00 00 .....
Transmitted SAPI:
00 00 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
  Policers: 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: Sendbroadcast-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:      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: 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: 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 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_Areal00
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      Bytes
                        (packet-per-second)  (bits-per-second)
      gold                  0              0
      (                    0) (              0)
      silver                0              0
      (                    0) (              0)

Logical interface so-0/1/3.0

      Destination class      Packets      Bytes
                        (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
                    iso
....

```

### show interfaces flow-statistics (Gigabit Ethernet)

user@host> show interfaces flow-statistics ge-0/0/1.0

```

Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 5161
  Output packets: 83
  Security: Zone: zone2
  Allowed host-inbound traffic : bootp bfd bgp  dns dvmrp  ldp msdp nhrp ospf
pgm
  pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
https ike
  netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
xnm-ssl
  lsping
  Flow Statistics :
  Flow Input statistics :
    Self packets :                0
    ICMP packets :                0
    VPN packets :                2564
    Bytes permitted by policy :    3478
    Connections established :     1
  Flow Output statistics:
    Multicast packets :           0
    Bytes permitted by policy :    16994
  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
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



|        |   |        |   |         |
|--------|---|--------|---|---------|
|        | ( | 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                                            |   |                     |   |                   |
| Destination class                                                         |   | Packets             |   | Bytes             |
|                                                                           |   | (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                                                         |   |                     |   |                   |



|                                                |                    |                         |                |            |
|------------------------------------------------|--------------------|-------------------------|----------------|------------|
| 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  |
| TCA-RAISED                                     |                        |           |                |            |
| BER                                            | 3.6e-5                 | 5.8e-5    | 3.6e-5         | 10.0e-3    |
| Yes                                            |                        |           |                | No         |
| Physical interface: et-0/1/0, SNMP ifIndex 515 |                        |           |                |            |
| 14:45-current                                  |                        |           |                |            |
| Suspect Flag:True                              | Reason:Object Disabled |           |                |            |
| PM                                             | CURRENT                | MIN       | MAX            | AVG        |
| 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            | 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 input power(0.01dBm)       | -3642 | -3665 | -3626 | -3637 | -1800 |
| -500 No No No No              |       |       |       |       |       |
| Module temperature(Celsius)   | 46    | 46    | 46    | 46    | -5    |
| 75 No No No No                |       |       |       |       |       |
| Tx laser bias current(0.1mA)  | 0     | 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 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 (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)

### List of Syntax

[Syntax \(M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface\) on page 1763](#)

[Syntax \(M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface\) on page 1763](#)

### 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>
```

### Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

### 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

### List of Sample Output

[show interfaces brief \(Management Ethernet\) on page 1768](#)

[show interfaces \(Management Ethernet\) on page 1768](#)

[show interfaces \(Management Ethernet \[TX Matrix Plus Router\]\) on page 1769](#)

[show interfaces \(Management Ethernet \[PTX Series Packet Transport Routers\]\) on page 1769](#)

[show interfaces detail \(Management Ethernet\) on page 1770](#)

[show interfaces detail \(Management Ethernet \[TX Matrix Plus Router\]\) on page 1771](#)

[show interfaces detail \(Management Ethernet \[PTX Packet Transport Routers\]\) on page 1772](#)

[show interfaces extensive \(Management Ethernet\) on page 1773](#)

[show interfaces extensive \(Management Ethernet \[TX Matrix Plus Router\]\) on page 1774](#)

[show interfaces extensive \(Management Ethernet \[PTX Series Packet Transport Routers\]\) on page 1775](#)

[show interfaces mgmtre0 \(Management Ethernet \[PTX5000 Router\]\) on page 1776](#)

[show interfaces brief \(Management Ethernet\) on page 1777](#)

[show interfaces brief \(Management Ethernet \[TX Matrix Plus Router\]\) on page 1777](#)

[show interfaces brief \(Management Ethernet \[PTX Series Packet Transport Routers\]\) on page 1777](#)

[show interfaces \(Internal Ethernet\) on page 1778](#)

[show interfaces \(Internal Ethernet \[TX Matrix Plus Router\]\) on page 1779](#)

[show interfaces detail \(Internal Ethernet\) on page 1779](#)

[show interfaces detail \(Internal Ethernet \[TX Matrix Plus Router\]\) on page 1780](#)

[show interfaces extensive \(internal Ethernet\) on page 1782](#)

[show interfaces extensive \(internal Ethernet \[TX Matrix Plus Router\]\) on page 1783](#)

### Output Fields

[Table 149 on page 1765](#) 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 149: 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 none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                      | <b>detail extensive none</b> |
| <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                   |
| <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 none</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>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>      |

Table 149: 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>Current address</b>         | Configured MAC address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail extensive none</b> |
| <b>Hardware address</b>        | Media access control (MAC) address of the interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive none</b> |
| <b>Alternate link address</b>  | Backup link address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <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 hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive none</b> |
| <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>      | <p>Number and rate of bytes and packets received and transmitted on the logical and physical interface.</p> <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>      |
| <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>Runs</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>             |

Table 149: 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>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 aborts 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> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                              |
| <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 none</b> |
| <b>SNMP ifIndex</b>      | Logical interface SNMP interface index number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive none</b> |
| <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 none</b> |
| <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 none</b> |
| <b>MTU</b>               | MTU size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</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 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 none</b> |



**Table 149: 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>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 none</b> |
| <b>Destination</b>      | IP address of the remote side of the connection.                                                                                                | <b>detail extensive none</b> |
| <b>Local</b>            | IP address of the logical interface.                                                                                                            | <b>detail extensive none</b> |
| <b>Broadcast</b>        | Broadcast address.                                                                                                                              | <b>detail extensive none</b> |
| <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 fxpl.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
```

## show interfaces (PPPoE)

### Syntax

```
show interfaces pp0.logical
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Release Information

Command introduced before Junos OS Release 7.4.

### Description

(M120 routers, M320 routers, and MX Series routers only). Display status information about the PPPoE interface.

### Options

**pp0.logical**—Display standard status information about the PPPoE interface.

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

**descriptions**—(Optional) Display interface description strings.

**media**—(Optional) Display media-specific information about PPPoE interfaces.

**snmp-index *snmp-index***—(Optional) Display information for the specified SNMP index of the interface.

**statistics**—(Optional) Display PPPoE interface statistics.

### Required Privilege Level

view

### List of Sample Output

[show interfaces \(PPPoE\) on page 1794](#)

[show interfaces \(PPPoE over Aggregated Ethernet\) on page 1794](#)

[show interfaces brief \(PPPoE\) on page 1795](#)

[show interfaces detail \(PPPoE\) on page 1795](#)

[show interfaces extensive \(PPPoE on M120 and M320 Routers\) on page 1796](#)

### Output Fields

[Table 150 on page 1787](#) lists the output fields for the **show interfaces (PPPoE)** command. Output fields are listed in the approximate order in which they appear.

Table 150: 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 <i>Common Output Fields Description</i> .               | All levels                   |
| <b>Interface index</b>    | Physical interface index number, which reflects its initialization sequence.                                                                       | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                      | <b>detail extensive none</b> |
| <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                   |
| <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>          | 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 <i>Common Output Fields Description</i> .         | 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>      |

Table 150: show interfaces (PPPoE) 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>        | MAC address of the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b> |
| <b>Alternate link address</b>  | Backup address of the link.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail extensive</b> |
| <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 150: 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 aborts 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>             |
| <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 aborts 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 none</b> |
| <b>SNMP ifIndex</b>      | Logical interface SNMP interface index number.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive none</b> |

Table 150: 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 <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | All levels              |
| <b>Encapsulation</b>      | Type of encapsulation configured on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | All levels              |
| <b>PPP parameters</b>     | PPP status: <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>              | PPPoE status: <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              |
| <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> |

Table 150: show interfaces (PPPoE) 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> <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>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.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> |
| <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 seconds</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-countnumber</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 number</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 150: 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 <b>hh:mm:ss</b>.</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 <b>hh:mm:ss</b>.</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>                                                                                                                                                                                                                                                                                                                                                                                                                                                   | none <b>detail extensive</b> |
| <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> none |

Table 150: show interfaces (PPPoE) Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Level of Output              |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <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> 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>MTU</b>                    | MTU size on the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> 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> none |
| <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> none |
| <b>Addresses, Flags</b>       | Information about the addresses configured for the protocol family. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> none |
| <b>Destination</b>            | IP address of the remote side of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail extensive</b> none |
| <b>Local</b>                  | IP address of the logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b> none |
| <b>Broadcast</b>              | Broadcast address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>detail extensive</b> 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: Sendbcst-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

```

## show interfaces interface-set (Ethernet Interface Set)

### Syntax

```
show interfaces interface-set interface-set-name
<detail | terse>
```

### Release Information

Command introduced in Junos OS Release 8.5.

### 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

## RELATED DOCUMENTATION

*Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

*Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

### List of Sample Output

[show interfaces interface-set terse on page 1801](#)

[show interfaces interface-set detail on page 1802](#)

[show interfaces interface-set \(ACI Interface Set based on ACI\) on page 1802](#)

[show interfaces interface-set \(ACI Interface Set based on ACI Trusted Option\) on page 1802](#)

[show interfaces interface-set \(ACI Interface Set based on ARI Trusted Option\) on page 1803](#)

[show interfaces interface-set \(ACI Interface Set based on ARI Trusted Option when both ACI and ARI are received\) on page 1803](#)

[show interfaces interface-set \(ACI Interface Set based on Accept-No-IDs Trusted Option when neither ACI nor ARI is received\) on page 1803](#)

[show interfaces interface-set \(L2BSA and PPPoE Subscribers\) on page 1804](#)

### Output Fields



Table 151 on page 1800 describes the information for the **show interfaces interface-set** command. Output fields are listed in the approximate order in which they appear.

**Table 151: 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> |
| <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> |

Table 151: Ethernet show interfaces interface-set Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                        | Level of Output    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| <b>PPPoE</b>                   | Dynamic PPPoE subscriber interface that the router creates using the ACI or ALI interface set.                                                                                                                                                                                                                                                                                                                                           | <b>detail</b> none |
| <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</b> none |
| <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</b> 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 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</b> none |

## 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

## show interfaces interface-set queue

### Syntax

```
show interfaces interface-set queue interface-set-name
<aggregate | remaining-traffic>
<forwarding-class class-name>
```

### Release Information

Command introduced in Junos OS Release 8.5.

### Description

Display information about the gigabit or 10-Gigabit Ethernet interface set queue. Supported in MX Series routers with enhanced queuing DPCs.

### Options

***interface-set-name***—(Optional) Display information about the specified gigabit or 10-Gigabit Ethernet interface set. Wildcard values can be used in the interface set name.

***aggregate***—(Optional) Display the aggregated queuing statistics of all member logical interfaces for interface sets that have traffic-control profiles configured.

***both-ingress-egress***—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics.

***egress***—(Optional) Display egress queue statistics.

***forwarding-class class-name***—(Optional) Display queuing statistics for the specified forwarding class.

***ingress***—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics.

***remaining-traffic***—(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

### List of Sample Output

[show interfaces interface-set queue \(Gigabit Ethernet\) on page 1807](#)

[show interfaces interface-set queue both-ingress-egress \(Enhanced DPC\) on page 1808](#)

[show interfaces interface-set queue egress \(Enhanced DPC\) on page 1811](#)

[show interfaces interface-set queue forwarding-class \(Gigabit Ethernet\) on page 1812](#)

[show interfaces interface-set queue \(Enhanced DPC\) on page 1813](#)

[show interfaces interface-set queue remaining-traffic \(Gigabit Ethernet\) on page 1814](#)

### Output Fields

Table 152 on page 1806 describes the information for the **show interfaces interface-set queue** command.

Table 152: 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      |
| <b>Queued</b>                       | <b>Packet</b> and <b>Byte</b> statistics for the specified queue. <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      |

Table 152: Ethernet show interfaces interface-set queue Output Fields (*continued*)

| Field Name         | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <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 |



# show interfaces interval

## Syntax

```
show interfaces interval
<interface-name>
```

## Release Information

Command introduced before Junos OS Release 7.4.

## 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

## RELATED DOCUMENTATION

[clear interfaces interval](#) | [1431](#)

## List of Sample Output

[show interfaces interval \(Channelized OC12\) on page 1817](#)

[show interfaces interval \(E3\) on page 1818](#)

[show interfaces interval \(SONET/SDH\) on page 1818](#)

## Output Fields

[Table 153 on page 1816](#) lists the output fields for the **show interfaces interval** command. Output fields are listed in the approximate order in which they appear.

Table 153: show interfaces interval Output Fields

| Field Name         | Field Description                             |
|--------------------|-----------------------------------------------|
| Physical interface | Name of the interface.                        |
| SNMP ifIndex       | SNMP index number for the physical interface. |

Table 153: show interfaces interval Output Fields (*continued*)

| Field Name                         | Field Description                                                                                                                                                                                                                                                                                                                                             |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>hh:mm-current</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>hh:mm-hh:mm</b>                 | Time of day (in hours and minutes) at the beginning and end of each 15-minute interval.                                                                                                                                                                                                                                                                       |
| <b>alarm or event: n</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>interface-type media</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><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

## show interfaces irb

### Syntax

```
show interfaces irb
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

### Release Information

Command introduced in Junos OS Release 8.4.

### Description

Display integrated routing and bridging interfaces information.

### Options

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

**descriptions**—(Optional) Display interface description strings.

**mac**—Display hardware MAC address

**media**—(Optional) Display media-specific information about network interfaces.

**snmp-index *snmp-index***—(Optional) Display information for the interface with the specified SNMP index.

**statistics**—(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

### List of Sample Output

[show interfaces irb extensive on page 1825](#)

[show interfaces irb snmp-index on page 1827](#)

### Output Fields

[Table 154 on page 1821](#) lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

Table 154: 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.                                                                                                              | terse                          |
| <b>Interface index</b>    | Physical interface index number, which reflects its initialization sequence.                                                                       | detail extensive<br>none       |
| <b>SNMP ifIndex</b>       | SNMP index number for the physical interface.                                                                                                      | detail extensive<br>none       |
| <b>Type</b>               | Physical interface type.                                                                                                                           | detail extensive<br>none       |
| <b>Link-level type</b>    | Encapsulation being used on the physical interface.                                                                                                | detail extensive<br>brief none |
| <b>MTU</b>                | MTU size on the physical interface.                                                                                                                | detail extensive<br>brief none |
| <b>Clocking</b>           | Reference clock source: <b>Internal</b> or <b>External</b> . Always unspecified on IRB interfaces.                                                 | detail extensive<br>brief      |
| <b>Speed</b>              | Speed at which the interface is running. Always unspecified on IRB interfaces.                                                                     | detail extensive<br>brief      |
| <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> . | detail extensive<br>brief none |
| <b>Interface flags</b>    | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .    | detail extensive<br>brief none |
| <b>Link type</b>          | Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .                                                                           | detail extensive<br>none       |
| <b>Link flags</b>         | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .             | detail extensive<br>none       |

Table 154: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Level of Output                 |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| <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><br>none |
| <b>Hardware address</b>        | MAC address of the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>detail extensive</b><br>none |
| <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.<br><br><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.<br><br><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 154: 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 aborts 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>         |
| <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 aborts 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 |



Table 154: show interfaces irb 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” 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>      | Number and rate of bytes and packets received and transmitted on the logical interface.<br><br><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 logical interface if IPv6 statistics tracking is enabled.<br><br><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>         |

Table 154: show interfaces irb 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>         |
| <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, 0 refers to the routing table inet.0.                                             | <b>detail extensive</b>         |
| <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 mac-database

### 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>
```

### 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.

### 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.

**ge-fpc/pic/port.n**—Display MAC addresses that have been learned on a particular logical interface.

**aex**—Display MAC addresses that have been learned on a particular aggregated Ethernet interface.

**xe-fpc/pic/port**—Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.

**xe-fpc/pic/port.n**—Display MAC addresses that have been learned on a particular logical interface.

**et-fpc/pic/port**—Display MAC addresses that have been learned on all logical interfaces on a particular physical interface.

**et-fpc/pic/port.n**—Display MAC addresses that have been learned on a particular logical interface.

**mac-address mac-address**—(Optional) Display detailed MAC address statistics, including policer information for **ge**, **xe**, and **et** 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

### List of Sample Output

[show interfaces mac-database \(All MAC Addresses on a Port\) on page 1831](#)

[show interfaces mac-database \(All MAC Addresses on an Aggregated Ethernet Interface\) on page 1832](#)

[show interfaces mac-database \(All MAC Addresses on a Service\) on page 1833](#)

[show interfaces mac-database mac-address on page 1833](#)

## Output Fields

Table 155 on page 1829 lists the output fields for the **show interfaces mac-database** command. Output fields are listed in the approximate order in which they appear.

Table 155: show interfaces mac-database 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 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 155: 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: <i>year-month-day hours:minutes:seconds timezone (wwEEKSddays hours:minutes ago)</i></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> ).                                                                                                                                                                 |
| <b>Encapsulation</b>                                                       | Encapsulation on the logical interface.                                                                                                                                                                                                                                                                                          |
| <b>MAC address, Input frames, Input bytes, Output frames, 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.                                                                                                                                                                                                                                                                                              |

Table 155: show interfaces mac-database Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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       |

## show interfaces mc-ae

### Syntax

```
show interfaces mc-ae
extensive
revertive-info
<id identifier unit number>
```

### 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

### 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

### RELATED DOCUMENTATION

*Configuring Multichassis Link Aggregation on MX Series Routers*

### List of Sample Output

[show interfaces mc-ae on page 1838](#)

[show interfaces mc-ae \(Active/Active Bridging and VRRP over IRB on MX Series Routers\) on page 1838](#)

[show interfaces mc-ae revertive-info on page 1839](#)

[show interfaces mc-ae extensive on page 1839](#)

[show interfaces mc-ae extensive \(MX Series Router after a configuration exchange error\) on page 1840](#)

[show interfaces mc-ae extensive on page 1840](#)

## Output Fields

Table 156 on page 1836 lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

Table 156: 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Peer State</b>                    | <p>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:</p> <p><b>Logical Interface</b>—Aggregated Ethernet (AE) aggregate number and unit number.</p> <p><b>Topology Type</b>—The bridge or VRRP topology type configured on the AE.</p> <p><b>Local State</b>—Up or down state of the local device.</p> <p><b>Peer State</b>—Up or down state of the peer device.</p> <p><b>Peer Ip/ICL-PL/State</b>—Address, interface and state of the peer device.</p> |
| <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> .                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 156: show interfaces mc-ae Output Fields (*continued*)

| Output Field Name                 | Field Description                                                                                                                                                                                                                        |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>switchover status</b>          | Status of the switchover if the <b>revert-time</b> statement is configured at the <b>[edit interfaces aex mc-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.                                                                                                                                                                                                      |
| 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.                                                                                                                                                                                                           |

Table 156: show interfaces mc-ae Output Fields (*continued*)

| Output Field Name              | Field Description                                                                  |
|--------------------------------|------------------------------------------------------------------------------------|
| <b>Admin Key</b>               | LACP administrative key of the node. Value should be the same on both MCLAG peers. |
| LACP Information               |                                                                                    |
| <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
```

## show interfaces prbs-stats

Syntax (MX10003 and MX204)

```
show interfaces interface-name prbs-stats
```

### Release Information

Statement introduced in Junos OS Release 19.2R1 for MX10003 and MX204 routers.

### 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](#) and [prbs-test-stop](#) 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 910](#).

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

## RELATED DOCUMENTATION

---

[prbs-test-start](#) | [1473](#)

---

[prbs-test-stop](#) | [1475](#)

---

[clear interfaces statistics](#)

---

[Verifying Link and Transceivers using Pseudo Random Binary Sequence \(PRBS\) Test](#) | [910](#)

### List of Sample Output

[show interfaces et-0/1/2 prbs-stats \(MX10003 and MX204 routers\) on page 1843](#)

## 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.

## show interfaces transport pm

### Syntax

```
show interfaces transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all | interface-name)
```

### Release Information

Command introduced in Junos OS Release 14.2 on the PTX Series.

Command introduced in Junos OS Release 16.1 on the MX Series.

Command introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

### 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

### RELATED DOCUMENTATION

[clear interfaces transport pm | 1435](#)

[tca | 972](#)

[transport-monitoring | 975](#)

### List of Sample Output

[show interfaces transport pm on page 1846](#)

[show interfaces transport \(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC\) on page 1848](#)

[show interfaces transport pm \(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC\) on page 1848](#)

[show interfaces transport \(PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC\) on page 1850](#)

[show interfaces transport pm optics \(PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC\) on page 1850](#)

[show interfaces transport pm otn \(PTX3000 router with 5-port 100-Gigabit DWDM OTN PIC\) on page 1851](#)  
[show interfaces transport pm optics \(MX10003, MX10008, MX10016, and MX204 router with QSFP-100GE-DWDM2 transceiver\) on page 1851](#)

### Output Fields

[Table 157 on page 1845](#) lists the output fields for the **show interfaces transport pm optics** command. Fields are listed in the approximate order in which they appear.

Table 157: show interfaces transport pm Output Fields

| Field Name         | Field Description                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------|
| Physical interface | Name of the physical interface.                                                                     |
| Interval           | The 15 minute interval for performance monitoring.                                                  |
| Suspect Flag       | TRUE if the performance monitoring data for the interval ap                                         |
| Reason             | Reason for setting the suspect flag.                                                                |
| COUNT              | Measured value.                                                                                     |
| THRESHOLD          | Threshold value set.                                                                                |
| TCA-ENABLED        | Threshold crossing alert. Set to TRUE if enabled.                                                   |
| TCA-RAISED         | TRUE if enabled and the value crosses the threshold.                                                |
| Near End PM        | Near end threshold crossing defect trigger. For more inform                                         |
| Far End PM         | Far end threshold crossing defect trigger. For more informat                                        |
| FEC PM             | Forwarding equivalence class threshold crossing defect trigge<br>trigger, see <a href="#">tca</a> . |
| BER PM             | Bit error rate threshold crossing defect trigger. For more info                                     |
| CURRENT            | Current value measured.                                                                             |
| PM                 | Performance monitor.                                                                                |
| MIN                | Minimum value measured.                                                                             |
| MAX                | Maximum value measured.                                                                             |

Table 157: show interfaces transport pm Output Fields (*continued*)

| Field Name                     | Field Description                                                                                 |
|--------------------------------|---------------------------------------------------------------------------------------------------|
| 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.                                                               |
| Module temperature (Celsius)   | The laser's temperature.                                                                          |
| Tx Laser bias current (0.1mA)  | Magnitude of the laser bias power setting current. The laser 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  |
| TCA-RAISED                                     |                        |           |                |            |
| BER                                            | 3.6e-5                 | 5.8e-5    | 3.6e-5         | 10.0e-3    |
| Yes                                            |                        |           |                | No         |
| Physical interface: et-0/1/0, SNMP ifIndex 515 |                        |           |                |            |
| 14:45-current                                  |                        |           |                |            |
| Suspect Flag:True                              | Reason:Object Disabled |           |                |            |
| PM                                             | CURRENT                | MIN       | MAX            | AVG        |
| 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   | 120   | 0     |
| 0                             | NA | NA | NA | NA    |       |       |       |       |
| SNR                           |    |    |    | 28    | 28    | 29    | 28    | 0     |
| 0                             | NA | NA | NA | NA    |       |       |       |       |
| Tx output power(0.01dBm)      |    |    |    | -5000 | -5000 | -5000 | -5000 | -300  |
| -100                          | No | No | No | No    |       |       |       |       |
| Rx output power(0.01dBm)      |    |    |    | -3642 | -3665 | -3626 | -3637 | -1800 |
| -500                          | No | No | No | No    |       |       |       |       |
| Module temperature(Celsius)   |    |    |    | 46    | 46    | 46    | 46    | -5    |
| 75                            | No | No | No | No    |       |       |       |       |
| Tx laser bias current(0.1mA)  |    |    |    | 0     | 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-RAISED           |                    |           |                       |            |
| BER                  | 4.0e-7             | 5.9e-7    | 5.1e-7                | 1.0e-2     |
| No                   |                    |           |                       | No         |

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)
(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 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)
(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      CURRENT      MIN      MAX      AVG      THRESHOLD

```

| TCA-ENABLED                  |       |       |       |         | TCA-RAISED |         |         |           |       |
|------------------------------|-------|-------|-------|---------|------------|---------|---------|-----------|-------|
|                              |       |       |       |         | (MIN)      |         |         |           |       |
| (MAX)                        | (MIN) | (MAX) | (MIN) | (MAX)   |            |         |         |           |       |
| 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                           |       |       |       | CURRENT | MIN        | MAX     | AVG     | THRESHOLD |       |
| TCA-ENABLED                  |       |       |       |         | TCA-RAISED |         |         |           |       |
|                              |       |       |       |         | (MIN)      |         |         |           |       |
| (MAX)                        | (MIN) | (MAX) | (MIN) | (MAX)   |            |         |         |           |       |
| 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                           |       |       |       | CURRENT | MIN        | MAX     | AVG     | THRESHOLD |       |
| TCA-ENABLED                  |       |       |       |         | TCA-RAISED |         |         |           |       |
|                              |       |       |       |         | (MIN)      |         |         |           |       |
| (MAX)                        | (MIN) | (MAX) | (MIN) | (MAX)   |            |         |         |           |       |
| 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 |     |     |     |       |

## show l2-learning instance

### Syntax

```
show l2-learning instance
```

### Release Information

(MX Series routers only) Command introduced in Junos OS Release 8.4.

### Description

Display Layer 2 learning properties for all the configured routing instances.

### Options

This command has no options.

### Required Privilege Level

view

### List of Sample Output

[show l2-learning instance on page 1855](#)

### Output Fields

[Table 158 on page 1854](#) describes the output fields for the **show l2-learning instance** command. Output fields are listed in the approximate order in which they appear.

Table 158: show l2-learning instance Output Fields

| Field Name       | Field Description                                                                                                                                                                                                                                   |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Routing Instance | Name of routing instance.                                                                                                                                                                                                                           |
| Bridging Domain  | 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. |
| Index            | Number associated with the routing instance or bridging domain.                                                                                                                                                                                     |
| Logical System   | Name of logical system or <b>Default</b> if no logical system is configured.                                                                                                                                                                        |

Table 158: 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 Instance     | Bridging Domain | Index | Logical System | Routing flags | MAC limit |
|----------------------|-----------------|-------|----------------|---------------|-----------|
| __juniper_private1__ |                 | 1     | Default        |               | 5000      |
| vs1                  | vlan100         | 3     | Default        |               | 5120      |
| vs1                  | vlan200         | 4     | Default        |               | 5120      |



## show l2-learning redundancy-groups

### Syntax

```
show l2-learning redundancy-groups
logical-system [system-name | all]
<redundancy-group-id [0 to 4294967294]>
arp-statistics
nd-statistics
remote-macs
```

### Release Information

Command introduced in Junos OS Release 13.2.

Support for logical systems added in Junos OS Release 14.1.

Command introduced in Junos OS Release 15.1R1 for EX Series switches

### 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.

**redundancy-group-id**—(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.

**arp-statistics**—(Optional) Count of ARP packets sent and received by the two MC-AE nodes.

**nd-statistics**—(Optional) Count of Neighbor Discovery packets sent and received by the two MC-AE nodes.

**remote-macs**—(Optional) List of remote MAC addresses in the “Installed” state, as learned from the remote MC-AE node.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[Configuring Multichassis Link Aggregation on MX Series Routers](#)

[show interfaces mc-ae | 1835](#)

*Configuring Active-Active Bridging and VRRP over IRB in Multichassis Link Aggregation on MX Series Routers and QFX Series Switches*

*Configuring Multichassis Link Aggregation on EX Series Switches*

### List of Sample Output

[show l2-learning redundancy-groups arp-statistics on page 1859](#)

[show l2-learning redundancy-groups nd-statistics on page 1859](#)

[show l2-learning redundancy-groups remote-macs on page 1859](#)

[show l2-learning redundancy-groups logical-system arp-statistics \(for Logical Systems\) on page 1860](#)

[show l2-learning redundancy-groups logical-system nd-statistics \(for Logical Systems\) on page 1860](#)

[show l2-learning redundancy-groups group-id on page 1860](#)

[show l2-learning redundancy-groups logical-system on page 1860](#)

### Output Fields

Output fields are listed in the approximate order in which they appear.

**Table 159: 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.                                   |
| 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 160: show l2-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 (MC-LAG) 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                 |
| Service-id                       | Service ID (configured at the routing instance level).                                        |

Table 161: show l2-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 l2-learning redundancy-groups arp-statistics

user@host> show l2-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

## show lacp interfaces

### Syntax

```
show lacp interfaces
<interface-name>
extensive
```

### Release Information

Command introduced in Junos OS Release 7.6.

**extensive** statement introduced in Junos OS Release 16.1R1

Command introduced in Junos OS Release 10.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Command introduced in Junos OS Release 14.2R3

### Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interface.

### Options

**none**—Display LACP information for all interfaces.

**interface-name**—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**aenumber**
- Fast Ethernet—**fe-fpc/pic/port**
- Gigabit Ethernet—**ge-fpc/pic/port**
- 10 Gigabit Ethernet—**xe-fpc/pic/port**

**extensive**—Display LACP information for the interface in detail.

**NOTE:** The **show lacp interfaces** command returns the following error message if your system is not configured in either active or passive LACP mode:

“Warning: lacp subsystem not running – not needed by configuration”

### Required Privilege Level

view

RELATED DOCUMENTATION

|                                                                                                 |
|-------------------------------------------------------------------------------------------------|
| <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure)</a>                           |
| <a href="#">Configuring Link Aggregation</a>                                                    |
| <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure)</a>                            |
| <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure)</a>                            |
| <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches</a> |
| <a href="#">Understanding Aggregated Ethernet Interfaces and LACP for Switches</a>              |
| <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a>                            |

List of Sample Output

[show lacp interfaces \(Aggregated Ethernet\) on page 1867](#)

[show lacp interfaces \(Gigabit Ethernet\) on page 1868](#)

[show lacp interfaces \(10 Gigabit Ethernet\) on page 1868](#)

Output Fields

[Table 162 on page 1863](#) lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 162: show lacp interfaces Output Fields

| Field Name           | Field Description                                                                                                    | Level of Out |
|----------------------|----------------------------------------------------------------------------------------------------------------------|--------------|
| LACP State           | For a child interface configured with the force-up statement, LACP state displays FUP along with the interface name. | All Levels   |
| Aggregated interface | Aggregated interface value.                                                                                          | All Levels   |



Table 162: show lacp interfaces Output Fields (continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Out |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| LACP State | <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>Long Timeout</b> or <b>Short Timeout</b>).</li> <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>• <b>Core isolation state down (CDN)</b>— LACP interface state. Down indicates the LACP interface is down because all the eBGP sessions for Ethernet VPN (EVPN) are down.</li> </ul> | All Levels   |

Table 162: show lacp interfaces Output Fields (continued)

| Field Name    | Field Description | Level of Out |
|---------------|-------------------|--------------|
| LACP Protocol |                   | All Levels   |

Table 162: show lacp interfaces Output Fields (continued)

| Field Name | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Out |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
|            | <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: <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> </ul> </li> </ul> |              |

Table 162: show lacp interfaces Output Fields (continued)

| Field Name       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Level of Out |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
|                  | <ul style="list-style-type: none"> <li>● <b>Waiting</b>—Multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |              |
| <b>LACP info</b> | <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           | priority      |        |     |                   | identifier |                         | priority |           | number   |
| key            |               |        |     |                   |            |                         |          |           |          |
| 1              | ge-0/0/1      | Actor  | 127 | 00:05:86:4e:b6:c0 |            |                         |          | 127       | 1        |

|   |          |         |     |                   |     |   |
|---|----------|---------|-----|-------------------|-----|---|
| 1 | ge-0/0/1 | Partner | 127 | 00:05:86:7e:d3:c0 | 127 | 1 |
| 1 | ge-0/0/2 | Actor   | 127 | 00:05:86:4e:b6:c0 | 127 | 2 |
| 1 | ge-0/0/2 | Partner | 127 | 00:05:86:7e:d3:c0 | 127 | 2 |
| 1 | ge-0/0/3 | Actor   | 127 | 00:05:86:4e:b6:c0 | 127 | 3 |
| 1 | ge-0/0/3 | Partner | 127 | 00:05:86:7e:d3:c0 | 127 | 3 |

### show lacp interfaces (Gigabit Ethernet)

user@host> show lacp interfaces ge-0/3/0

```

Aggregated interface: ae0
  LACP State:      Role   Exp   Def   Dist   Col   Syn   Aggr   Timeout   Activity
    ge-0/3/0      Actor  No    No    Yes   Yes   Yes   Yes     Fast     Active
    ge-0/3/0      Partner No    No    Yes   Yes   Yes   Yes     Fast     Active
  LACP Protocol:  Receive State   Transmit State       Mux State
    ge-0/3/0      Current      Fast periodic   Collecting distributing

```

### show lacp interfaces (10 Gigabit Ethernet)

user@host> show lacp interfaces xe-1/0/2

```

Aggregated interface: ae0
  LACP State:      Role   Exp   Def   Dist   Col   Syn   Aggr   Timeout   Activity
    xe-1/0/2      Actor  No    No    Yes   Yes   Yes   Yes     Fast     Active
    xe-1/0/2      Partner No    No    Yes   Yes   Yes   Yes     Fast     Active
  LACP Protocol:  Receive State   Transmit State       Mux State
    xe-1/0/2      Current      Fast periodic   Collecting distributing

```

## show oam ethernet connectivity-fault-management delay-statistics

### Syntax

```
show oam ethernet connectivity-fault-management delay-statistics
<count entry-count>
<local-mep mep-id>
<maintenance-association ma-name>
<maintenance-domain md-name>
<remote-mep remote-mep-id>
```

### Release Information

Command introduced in Junos OS Release 9.5.

Command introduced in Junos OS Release 11.4 for EX Series switches.

### Description

On MX Series routers with Ethernet interfaces on Dense Port Concentrators (DPCs), display ETH-DM delay statistics.

On EX Series switches, display delay measurement results.

### Options

**count entry-count**—(Optional) Number of entries to display from the statistics table. The range of values is **1** through **100**. The default value is **100** entries.

**local-mep mep-id**—(Optional) Numeric identifier of the local MEP. On MX Series routers, the range of values is **1** through **8192**. On EX Series switches, the range of values is **1** through **8191**.

**maintenance-association ma-name**—Name of an existing CFM maintenance association.

**maintenance-domain md-name**—Name of an existing connectivity fault management (CFM) maintenance domain.

**remote-mep remote-mep-id**—(Optional) Numeric identifier of the remote MEP. On MX Series routers, the range of values is **1** through **8192**. On EX Series switches, the range of values is **1** through **8191**.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[clear oam ethernet connectivity-fault-management statistics | 1444](#)

[clear oam ethernet connectivity-fault-management delay-statistics](#)

[show oam ethernet connectivity-fault-management interfaces | 1879](#)

[show oam ethernet connectivity-fault-management mep-database | 1901](#)

[show oam ethernet connectivity-fault-management mep-statistics | 1916](#)

### List of Sample Output

[show oam ethernet connectivity-fault-management delay-statistics on page 1871](#)

[show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 1872](#)

### Output Fields

Table 163 on page 1870 lists the output fields for the **show oam ethernet connectivity-fault-management delay-statistics** command and the **show oam ethernet connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 163: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields

| Output Field Name     | Field Description                                                                                                                                                                                                                                                                                                        |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MEP identifier        | Maintenance association end point (MEP) numeric identifier.                                                                                                                                                                                                                                                              |
| MAC address           | Unicast MAC address configured for the MEP.                                                                                                                                                                                                                                                                              |
| Remote MEP count      | Number of remote MEPs (unless you specify the <b>remote-mep</b> option).                                                                                                                                                                                                                                                 |
| Remote MEP identifier | Numeric identifier of the remote MEP.                                                                                                                                                                                                                                                                                    |
| Remote MAC address    | Unicast MAC address of the remote MEP.                                                                                                                                                                                                                                                                                   |
| Index                 | Index number that corresponds to the ETH-DM entry in the CFM database.                                                                                                                                                                                                                                                   |
| One-way delay (usec)  | <p>For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP.</p> <p>For a detailed description of one-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>  |
| Two-way delay (usec)  | <p>For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP.</p> <p>For a detailed description of two-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p> |
| Average one-way delay | Average one-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                |

Table 163: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields (continued)

| Output Field Name               | Field Description                                            |
|---------------------------------|--------------------------------------------------------------|
| Average one-way delay variation | Average one-way "frame jitter" for the statistics displayed. |
| Best-case one-way delay         | Lowest one-way frame delay for the statistics displayed.     |
| Worst-case one-way delay        | Highest one-way frame delay for the statistics displayed.    |
| Average two-way delay           | Average two-way frame delay for the statistics displayed.    |
| Average two-way delay variation | Average two-way "frame jitter" for the statistics displayed. |
| Best-case two-way delay         | Lowest two-way frame delay for the statistics displayed.     |
| Worst-case two-way delay        | Highest two-way frame delay calculated in this session.      |

## Sample Output

show oam ethernet connectivity-fault-  
management  
delay-statistics

user@switch> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain  
md6 maintenance-association ma6

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)         (usec)
  1      259          519
  2      273          550
  3      287          571
  4      299          610
  5      313          650
Average one-way delay           : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay         : 259 usec
```



```

Worst case one-way delay      : 313 usec
Average two-way delay        : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay      : 519 usec
Worst case two-way delay     : 650 usec

```

Remote MEP identifier: 102

Remote MAC address: 00:04:55:63:39:5a

Delay measurement statistics:

| Index | One-way delay<br>(usec) | Two-way delay<br>(usec) |
|-------|-------------------------|-------------------------|
| 1     | 29                      | 58                      |
| 2     | 23                      | 59                      |
| 3     | 27                      | 56                      |
| 4     | 29                      | 62                      |
| 5     | 33                      | 68                      |

```

Average one-way delay      : 28 usec
Average one-way delay variation: 3 usec
Best case one-way delay    : 23 usec
Worst case one-way delay   : 33 usec
Average two-way delay      : 60 usec
Average two-way delay variation: 3 usec
Best case two-way delay    : 56 usec
Worst case two-way delay   : 68 usec

```

### **show oam ethernet connectivity-fault-management delay-statistics remote-mep**

**user@switch> show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md6 maintenance-association ma6 remote-mep 101**

MEP identifier: 100, MAC address: 00:05:85:73:7b:39

Remote MEP identifier: 101

Remote MAC address: 00:05:85:73:39:4a

Delay measurement statistics:

| Index | One-way delay<br>(usec) | Two-way delay<br>(usec) |
|-------|-------------------------|-------------------------|
| 1     | 259                     | 519                     |
| 2     | 273                     | 550                     |
| 3     | 287                     | 571                     |
| 4     | 299                     | 610                     |
| 5     | 313                     | 650                     |

```

Average one-way delay      : 286 usec

```

```
Average one-way delay variation: 62 usec
Best case one-way delay          : 259 usec
Worst case one-way delay         : 313 usec
Average two-way delay            : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay          : 519 usec
Worst case two-way delay         : 650 usec
```

# show oam ethernet connectivity-fault-management forwarding-state

## Syntax

```
show oam ethernet connectivity-fault-management forwarding-state
interface interface-name | instance instance-name
<brief | detail | extensive>
```

## Release Information

Command introduced in Junos OS Release 8.4.

## Description

On M7i and M10i with the Enhanced CFEB (CFEB-E), M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.

## Options

**interface *interface-name***—Display forwarding state information for the specified Ethernet interface only.

**instance *instance-name***—Display forwarding state information for the specified forwarding instance only.

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

## Required Privilege Level

view

## List of Sample Output

- [show oam ethernet connectivity-fault-management forwarding-state instance on page 1875](#)
- [show oam ethernet connectivity-fault-management forwarding-state interface on page 1876](#)
- [show oam ethernet connectivity-fault-management forwarding-state interface detail on page 1876](#)
- [show oam ethernet connectivity-fault-management forwarding-state interfaceinterface-name on page 1878](#)

## Output Fields

[Table 164 on page 1874](#) lists the output fields for the **show oam ethernet connectivity-fault-management forwarding-state** command. Output fields are listed in the approximate order in which they appear.

Table 164: show oam ethernet connectivity-fault-management forwarding-state Output Fields

| Field Name     | Field Description                        | Level of Output |
|----------------|------------------------------------------|-----------------|
| Interface name | Interface identifier.                    | All levels      |
| Link (Status)  | Local link status.                       | All levels      |
| Filter action  | Filter action for messages at the level. | All levels      |

Table 164: show oam ethernet connectivity-fault-management forwarding-state Output Fields (*continued*)

| Field Name           | Field Description                        | Level of Output |
|----------------------|------------------------------------------|-----------------|
| <b>Next hop type</b> | Next-hop type.                           | All levels      |
| <b>Next index</b>    | Next-hop index number.                   | <b>brief</b>    |
| <b>Level</b>         | Maintenance domain (MD) level.           | <b>detail</b>   |
| <b>Direction</b>     | MEP direction configured.                | none            |
| <b>Instance name</b> | Forwarding instance name.                | All levels      |
| <b>CEs</b>           | Number of customer edge (CE) interfaces. | All levels      |
| <b>VEs</b>           | Number of VPN endpoint (VE) interfaces.  | All levels      |

## Sample Output

```
show oam ethernet
connectivity-fault-
management forwarding-
state instance
```

```
user@host> show oam ethernet connectivity-fault-management forwarding-state instance
```

```
Instance name: __+bd1__
CEs: 3
VEs: 0
Maintenance domain forwarding state:
```

| Level | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
|-------|-----------|---------------|-----------------|------------------|
| 0     |           | Drop          | none            |                  |
| 1     |           | Drop          | none            |                  |
| 2     |           | Drop          | none            |                  |
| 3     |           | Drop          | none            |                  |
| 4     |           | Drop          | none            |                  |
| 5     |           | Drop          | none            |                  |
| 6     |           | Drop          | none            |                  |
| 7     |           | Drop          | none            |                  |

**show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface**

user@host> **show oam ethernet connectivity-fault-management forwarding-state interface**

Interface name: ge-3/0/0.0

Instance name: \_\_+bd1\_\_

Maintenance domain forwarding state:

| Level | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
|-------|-----------|---------------|-----------------|------------------|
| 0     |           | Drop          | none            |                  |
| 1     |           | Drop          | none            |                  |
| 2     |           | Drop          | none            |                  |
| 3     |           | Drop          | none            |                  |
| 4     |           | Drop          | none            |                  |
| 5     |           | Drop          | none            |                  |
| 6     |           | Drop          | none            |                  |
| 7     | down      | Receive       | none            |                  |

Interface name: xe-0/0/0.0

Instance name: \_\_+bd1\_\_

Maintenance domain forwarding state:

| Level | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
|-------|-----------|---------------|-----------------|------------------|
| 0     |           | Drop          | none            |                  |
| 1     |           | Drop          | none            |                  |
| 2     |           | Drop          | none            |                  |
| 3     |           | Drop          | none            |                  |
| 4     |           | Drop          | none            |                  |
| 5     |           | Drop          | none            |                  |
| 6     |           | Drop          | none            |                  |
| 7     | down      | Receive       | none            |                  |

**show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface detail**

user@host> **show oam ethernet connectivity-fault-management forwarding-state interface detail**

Interface name: ge-3/0/0.0

Instance name: \_\_+bd1\_\_

Level: 0  
Filter action: Drop  
Nexthop type: none

Level: 1  
Filter action: Drop  
Nexthop type: none

Level: 2  
Filter action: Drop  
Nexthop type: none

Level: 3  
Filter action: Drop  
Nexthop type: none

Level: 4  
Filter action: Drop  
Nexthop type: none

Level: 5  
Filter action: Drop  
Nexthop type: none

Level: 6  
Filter action: Drop  
Nexthop type: none

Level: 7  
Direction: down  
Filter action: Receive  
Nexthop type: none

Interface name: xe-0/0/0.0  
Instance name: \_\_+bd1\_\_

Level: 0  
Filter action: Drop  
Nexthop type: none

Level: 1  
Filter action: Drop  
Nexthop type: none

...

show oam ethernet  
connectivity-fault-  
management forwarding-  
state interface  
interface-name

user@host> show oam ethernet connectivity-fault-management forwarding-state interface  
interface-name ge-3/0/0/0.0

|                                      |           |               |                 |                  |
|--------------------------------------|-----------|---------------|-----------------|------------------|
| Interface name: ge-3/0/0.0           |           |               |                 |                  |
| Instance name: __+bd1__              |           |               |                 |                  |
| Maintenance domain forwarding state: |           |               |                 |                  |
| Level                                | Direction | Filter action | Nexthop<br>type | Nexthop<br>index |
| 0                                    |           | Drop          | none            |                  |
| 1                                    |           | Drop          | none            |                  |
| 2                                    |           | Drop          | none            |                  |
| 3                                    |           | Drop          | none            |                  |
| 4                                    |           | Drop          | none            |                  |
| 5                                    |           | Drop          | none            |                  |
| 6                                    |           | Drop          | none            |                  |
| 7                                    | down      | Receive       | none            |                  |

## show oam ethernet connectivity-fault-management interfaces

### Syntax

```
show oam ethernet connectivity-fault-management interfaces
<ethernet-interface-name>
<level md-level>
<brief | detail | extensive>
```

### Release Information

Command introduced in Junos OS Release 8.4.

Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.

Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for ACX Series and MX Series routers.

### Description

On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.

In addition, for Ethernet interfaces on MX Series routers, also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts when **detail** or **extensive** mode is specified.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

### Options

**brief | detail | extensive**—(Optional) Specified level of output.

**ethernet-interface-name**—(Optional) CFM information only for CFM entities attached to the specified Ethernet interface.

**level md-level**—(Optional) CFM information for CFM identities enclosed within a maintenance domain of the specified level.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[clear oam ethernet connectivity-fault-management statistics](#) | 1444

[show oam ethernet connectivity-fault-management delay-statistics](#) | 1869

[show oam ethernet connectivity-fault-management mep-database](#) | 1901



[show oam ethernet connectivity-fault-management mep-statistics](#) | 1916

[Ethernet Alarm Indication Signal \(ETH-AIS\) Function Overview](#) | 810

### List of Sample Output

[show oam ethernet connectivity-fault-management interfaces](#) on page 1886

[show oam ethernet connectivity-fault-management interfaces detail](#) on page 1886

[show oam ethernet connectivity-fault-management interfaces detail \(One-Way ETH-DM\)](#) on page 1887

[show oam ethernet connectivity-fault-management interfaces detail\(Connection Protection TLV Configured\)](#) on page 1888

[show oam ethernet connectivity-fault-management interfaces extensive](#) on page 1889

[show oam ethernet connectivity-fault-management interfaces level](#) on page 1891

[show oam ethernet connectivity-fault-management interfaces \(trunk ports\)](#) on page 1891

### Output Fields

Table 165 on page 1880 lists the output fields for the **show oam ethernet connectivity-fault-management interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 165: show oam ethernet connectivity-fault-management interfaces Output Fields

| Field Name                              | Field Description                                                 | Level of Output         |
|-----------------------------------------|-------------------------------------------------------------------|-------------------------|
| <b>Interface</b>                        | Interface identifier.                                             | All levels              |
| <b>Interface status</b>                 | Local interface status.                                           | All levels              |
| <b>Link status</b>                      | Local link status. <b>Up</b> , <b>down</b> , or <b>oam-down</b> . | All levels              |
| <b>Maintenance domain name</b>          | Maintenance domain name.                                          | <b>detail extensive</b> |
| <b>Format (Maintenance domain)</b>      | Maintenance domain name format configured.                        | <b>detail extensive</b> |
| <b>Level</b>                            | Maintenance domain level configured.                              | All levels              |
| <b>Maintenance association name</b>     | Maintenance association name.                                     | <b>detail extensive</b> |
| <b>Format (Maintenance association)</b> | Maintenance association name format configured.                   | <b>detail extensive</b> |
| <b>Continuity-check status</b>          | Continuity-check status.                                          | <b>detail extensive</b> |

Table 165: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)

| Field Name                | Field Description                                                                                                                                                                                                                                                       | Level of Output         |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| Ethernet-ais status       | Status of alarm indication signal (AIS). <b>active</b> or <b>in-active</b> .                                                                                                                                                                                            | <b>detail extensive</b> |
| Interval                  | Continuity-check message interval.                                                                                                                                                                                                                                      | <b>detail extensive</b> |
| Loss-threshold            | Lost continuity-check message threshold.                                                                                                                                                                                                                                | <b>detail extensive</b> |
| Interface status TLV      | Status of the interface status TLV, if configured on the MEP interface: <b>none, up, down, testing, unknown, dormant, notPresent, lowerLayerDown</b>                                                                                                                    | <b>detail extensive</b> |
| Port status TLV           | Status of the port status TLV, if configured on the MEP interface: <b>none, no, yes</b>                                                                                                                                                                                 | <b>detail extensive</b> |
| Connection Protection TLV | Status of the connection protection TLV if configured on the MEP interface: <b>no, yes</b><br><br>If <b>yes</b> , then the transmitted connection protection TLV is decoded and the following three fields are displayed: <b>Prefer me, Protection in use, FRR Flag</b> | <b>detail extensive</b> |
| Prefer me                 | If set to <b>yes</b> , the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to the remote side.<br><br>Its value can be <b>yes</b> or <b>no</b> .                                          | <b>detail extensive</b> |
| Protection in use         | Used for protection decision coordination. Its value is set to <b>yes</b> if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path.<br><br>Its value can be <b>yes</b> or <b>no</b> .                                         | <b>detail extensive</b> |
| FRR Flag                  | LSR/LER forwarding the CCM Frame into a bypass tunnel is set.<br><br>Its value can be <b>yes</b> or <b>no</b> .                                                                                                                                                         | <b>detail extensive</b> |
| MEP identifier            | Maintenance association end point (MEP) identifier.                                                                                                                                                                                                                     | All levels              |
| Neighbors                 | Number of MEP neighbors.                                                                                                                                                                                                                                                | All levels              |
| Direction                 | MEP direction configured.                                                                                                                                                                                                                                               | <b>detail extensive</b> |
| MAC address               | MAC address configured for the MEP.                                                                                                                                                                                                                                     | <b>detail extensive</b> |

Table 165: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)

| Field Name                               | Field Description                                                                                                                                                                                              | Level of Output         |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>MEP status</b>                        | Indicates the status of the connectivity fault management (CFM) protocol running on the MEP: <b>Running, inactive, disabled, or unsupported.</b>                                                               | <b>detail extensive</b> |
| <b>Remote MEP not receiving CCM</b>      | Whether the remote MEP is not receiving connectivity check messages (CCMs).                                                                                                                                    | <b>detail extensive</b> |
| <b>Erroneous CCM received</b>            | Whether erroneous CCMs have been received.                                                                                                                                                                     | <b>detail extensive</b> |
| <b>Cross-connect CCM received</b>        | Whether cross-connect CCMs have been received.                                                                                                                                                                 | <b>detail extensive</b> |
| <b>RDI sent by some MEP</b>              | Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs. | <b>detail extensive</b> |
| Some remote MEP's MAC in error state     | Indicates whether the remote MEP's MAC is in error state.                                                                                                                                                      | <b>detail extensive</b> |
| Alarm Indication Signal                  | Indicates whether the AIS is triggered or is cleared.                                                                                                                                                          | <b>detail extensive</b> |
| <b>CCMs sent</b>                         | Number of CCMs transmitted.                                                                                                                                                                                    | <b>detail extensive</b> |
| <b>CCMs received out of sequence</b>     | Number of CCMs received out of sequence.                                                                                                                                                                       | <b>detail extensive</b> |
| <b>LBM sent</b>                          | Number of loopback request messages (LBMs) sent.                                                                                                                                                               | <b>detail extensive</b> |
| <b>Valid in-order LBRs received</b>      | Number of loopback response messages (LBRs) received that were valid messages and in sequence.                                                                                                                 | <b>detail extensive</b> |
| <b>Valid out-of-order LBRs received</b>  | Number of LBRs received that were valid messages and not in sequence.                                                                                                                                          | <b>detail extensive</b> |
| <b>LBRs received with corrupted data</b> | Number of LBRs received that were corrupted.                                                                                                                                                                   | <b>detail extensive</b> |

Table 165: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                                            | Level of Output         |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>LBRs sent</b>                           | Number of LBRs transmitted.                                                                                                                                                                                                  | <b>detail extensive</b> |
| <b>LTM sent</b>                            | Linktrace messages (LTMs) transmitted.                                                                                                                                                                                       | <b>detail extensive</b> |
| <b>LTM received</b>                        | Linktrace messages received.                                                                                                                                                                                                 | <b>detail extensive</b> |
| <b>LTR sent</b>                            | Linktrace responses (LTRs) transmitted.                                                                                                                                                                                      | <b>detail extensive</b> |
| <b>LTR received</b>                        | Linktrace responses received.                                                                                                                                                                                                | <b>detail extensive</b> |
| <b>Sequence number of next LTM request</b> | Sequence number of next LTM request to be transmitted.                                                                                                                                                                       | <b>detail extensive</b> |
| <b>1DMs sent</b>                           | If the interface is attached to an initiator MEP for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.<br><br>For all other cases, this field displays 0. | <b>detail extensive</b> |
| <b>Valid 1DMs received</b>                 | If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of valid 1DM frames received.<br><br>For all other cases, this field displays 0.                                                         | <b>detail extensive</b> |
| <b>Invalid 1DMs received</b>               | If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of invalid 1DM frames received.<br><br>For all other cases, this field displays 0.                                                       | <b>detail extensive</b> |
| <b>Out of sync 1DMs received</b>           | If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.                                                                       | <b>detail extensive</b> |
| <b>DMMs sent</b>                           | If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.<br><br>For all other cases, this field displays 0. | <b>detail extensive</b> |
| <b>Valid DMMs received</b>                 | If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid two-way delay measurement request packets received.                                                                           | <b>detail extensive</b> |

Table 165: show oam ethernet connectivity-fault-management interfaces Output Fields (continued)

| Field Name                          | Field Description                                                                                                                                                                     | Level of Output         |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>Invalid DMMs received</b>        | If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid two-way delay measurement request packets received.                                  | <b>detail extensive</b> |
| <b>DMRs sent</b>                    | If the interface is attached to a responder MEP for a two-way ETH-DM session: Number of delay measurement reply (DMR) frames sent.<br><br>For all other cases, this field displays 0. | <b>detail extensive</b> |
| <b>Valid DMRs received</b>          | If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid DMRs received.<br><br>For all other cases, this field displays 0.                      | <b>detail extensive</b> |
| <b>Invalid DMRs received</b>        | If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid DMRs received.<br><br>For all other cases, this field displays 0.                    | <b>detail extensive</b> |
| <b>LMM sent</b>                     | If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.                      | <b>detail extensive</b> |
| <b>Valid LMM received</b>           | If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.                                                     | <b>detail extensive</b> |
| <b>Invalid LMM received</b>         | If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid loss measurement request packets received.                                                   | <b>detail extensive</b> |
| <b>LMR sent</b>                     | If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.                                                             | <b>detail extensive</b> |
| <b>Valid LMR received</b>           | If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid LMR frames received.                                                                           | <b>detail extensive</b> |
| <b>Invalid LMR received</b>         | If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.                                                                         | <b>detail extensive</b> |
| <b>Valid AIS frames transmitted</b> | Number of valid AIS frames transmitted to the peer MEPs.                                                                                                                              | <b>detail extensive</b> |
| <b>Valid AIS frames received</b>    | Number of valid AIS frames received from the peer MEPs.                                                                                                                               | <b>detail extensive</b> |

Table 165: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                                                                                       | Level of Output         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <b>SLM sent</b>                 | If the interface is attached to an initiator MEP for a ETH-SLM session:<br>Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session. | <b>detail extensive</b> |
| <b>Valid SLM received</b>       | If the interface is attached to a responder MEP for a ETH-SLM session:<br>Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.                                                    | <b>detail extensive</b> |
| <b>Invalid SLM received</b>     | If the interface is attached to a responder MEP for a ETH-SLM session:<br>Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.                                                  | <b>detail extensive</b> |
| <b>SLR sent</b>                 | If the interface is attached to a responder MEP for a ETH-SLM session:<br>Number detail extensive of synthetic loss reply (SLR) frames sent.                                                                            | <b>detail extensive</b> |
| <b>Valid SLR received</b>       | If the interface is attached to an initiator MEP for a ETH-SLM session:<br>Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.                                                    | <b>detail extensive</b> |
| <b>Invalid SLR received</b>     | If the interface is attached to an initiator MEP for a ETH-SLM session:<br>Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.                                                  | <b>detail extensive</b> |
| <b>Remote MEP count</b>         | Number of remote MEPs.                                                                                                                                                                                                  | <b>extensive</b>        |
| <b>Identifier (remote MEP)</b>  | MEP identifier of the remote MEP.                                                                                                                                                                                       | <b>extensive</b>        |
| <b>MAC address (remote MEP)</b> | MAC address of the remote MEP.                                                                                                                                                                                          | <b>extensive</b>        |
| <b>State (remote MEP)</b>       | State of the remote MEP.                                                                                                                                                                                                | <b>extensive</b>        |
| <b>Interface (remote MEP)</b>   | Interface of the remote MEP.                                                                                                                                                                                            | <b>extensive</b>        |

## Sample Output

**show oam ethernet connectivity-fault-management interfaces**

user@host> **show oam ethernet connectivity-fault-management interfaces**

| Interface    | Link | Status | Level | MEP Identifier | Neighbors |
|--------------|------|--------|-------|----------------|-----------|
| ge-1/1/0.0   | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.1   | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.10  | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.100 | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.101 | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.102 | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.103 | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.104 | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.105 | Up   | Active | 0     | 2              | 1         |
| ge-1/1/0.106 | Up   | Active | 0     | 2              | 1         |
| ...          |      |        |       |                |           |

**show oam ethernet connectivity-fault-management interfaces detail**

user@host> **show oam ethernet connectivity-fault-management interfaces detail**

```

Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : yes
  Cross-connect CCM received                   : no
  RDI sent by some MEP                        : yes
  Alarm Indication Signal                     : yes
Statistics:
  CCMs sent                                   : 76
  CCMs received out of sequence               : 0
  LBMs sent                                   : 0
  Valid in-order LBRs received                : 0
  Valid out-of-order LBRs received            : 0

```

```

LBRs received with corrupted data      : 0
LBRs sent                              : 0
LTMs sent                              : 0
LTMs received                          : 0
LTRs sent                              : 0
LTRs received                          : 0
Sequence number of next LTM request    : 1320235363

LDMs sent                              : 0
Valid LDMs received                    : 0
Invalid LDMs received                  : 0
DMMs sent                              : 0
DMRs sent                              : 0
Valid DMRs received                    : 0
Invalid DMRs received                  : 0
LMM sent                               : 10
Valid LMM received                     : 20
Invalid LMM received                   : 0
LMR sent                               : 20
Valid LMR received                     : 10
Invalid LMR received                   : 0
Valid AIS frames transmitted           : 0
Valid AIS frames received               : 0
SLM sent                               : 10
Valid SLM received                     : 20
Invalid SLM received                   : 0
SLR sent                               : 20
Valid SLR received                     : 10
Invalid SLR received                   : 0
Remote MEP count: 2
Identifier    MAC address    State    Interface
2001         00:90:69:0b:7f:71  ok      ge-5/2/9.0
4001         00:90:69:0b:09:c5  ok      ge-5/2/9.0

```

**show oam ethernet connectivity-fault-management interfaces detail (One-Way ETH-DM)**

user@host **show oam ethernet connectivity-fault-management interfaces detail**

```

Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57

```



```

MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMs sent                                  : 1590
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                             : 0
  LTRs sent                                  : 0
  LTRs received                             : 0
  Sequence number of next LTM request        : 1542035464

  1DMs sent                                  : 10
  Valid 1DMs received                       : 0
  Invalid 1DMs received                     : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                      : 0
  Invalid DMRs received                    : 0
Remote MEP count: 1
  Identifier   MAC address   State   Interface
    201       00:90:69:0a:43:94   ok     ge-0/2/5.0

```

**show oam ethernet connectivity-fault-management interfaces detail  
(Connection Protection TLV Configured)**

user@host#show oam ethernet connectivity-fault-management interfaces detail

```

Interface name: xe-6/2/0.0 , Interface status: Active, Link status: Up
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: yes

```

```

    Prefer me: no, Protection in use: no, FRR Flag: no
MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
MEP status: running
Defects:
    Remote MEP not receiving CCM                : no
    Erroneous CCM received                       : no
    Cross-connect CCM received                   : no
    RDI sent by some MEP                        : no
    Some remote MEP's MAC in error state         : no
Statistics:
    CCMs sent                                   : 225
    CCMs received out of sequence                : 0
    LBMs sent                                   : 0
    Valid in-order LBRs received                 : 0
    Valid out-of-order LBRs received             : 0
    LBRs received with corrupted data            : 0
    LBRs sent                                   : 0
    LTMs sent                                   : 0
    LTMs received                               : 0
    LTRs sent                                   : 0
    LTRs received                               : 0
    Sequence number of next LTM request          : 1244305646

    1DMs sent                                   : 0
    Valid 1DMs received                         : 0
    Invalid 1DMs received                       : 0
    Out of sync 1DMs received                   : 0
    DMMS sent                                   : 0
    Valid DMMS received                         : 0
    Invalid DMMS received                       : 0
    DMRs sent                                   : 0
    Valid DMRs received                         : 0
    Invalid DMRs received                       : 0
    LMMS sent                                   : 0
    Valid LMMS received                         : 0
    Invalid LMMS received                       : 0
    LMRs sent                                   : 0
    Valid LMRs received                         : 0
    Invalid LMRs received                       : 0
Remote MEP count: 1
    Identifier    MAC address    State    Interface
      2          00:90:69:7f:e4:30

```

**show oam ethernet connectivity-fault-  
management interfaces**

**extensive**

user@host> **show oam ethernet connectivity-fault-management interfaces extensive**

```

Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
    Remote MEP not receiving CCM                : no
    Erroneous CCM received                       : yes
    Cross-connect CCM received                   : no
    RDI sent by some MEP                        : yes

    Alarm Indication Signal                     : yes
Statistics:
    CCMs sent                                   : 76
    CCMs received out of sequence                : 0
    LBMs sent                                   : 0
    Valid in-order LBRs received                 : 0
    Valid out-of-order LBRs received             : 0
    LBRs received with corrupted data            : 0
    LBRs sent                                   : 0
    LTMs sent                                   : 0
    LTMs received                               : 0
    LTRs sent                                   : 0
    LTRs received                               : 0
    Sequence number of next LTM request          : 1645032434

    1DMs sent                                   : 0
    Valid 1DMs received                         : 0
    Invalid 1DMs received                       : 0
    DMMs sent                                   : 0
    DMRs sent                                   : 0
    Valid DMRs received                         : 0
    Invalid DMRs received                       : 0

    Valid AIS frames transmitted                 : 0
    Valid AIS frames received                   : 0
    SLM sent                                    : 10
    Valid SLM received                           : 20

```

```

Invalid SLM received           : 0
SLR sent                       : 20
Valid SLR received             : 10
Invalid SLR received           : 0
Remote MEP count: 2
Identifier    MAC address      State    Interface
2001         00:90:69:0b:7f:71  ok      ge-5/2/9.0
4001         00:90:69:0b:09:c5  ok      ge-5/2/9.0

```

#### show oam ethernet connectivity-fault-management interfaces level

user@host> show oam ethernet connectivity-fault-management interfaces level 7

| Interface  | Link | Status | Level | MEP Identifier | Neighbors |
|------------|------|--------|-------|----------------|-----------|
| ge-3/0/0.0 | Up   | Active | 7     | 201            | 0         |
| xe-0/0/0.0 | Up   | Active | 7     | 203            | 1         |

#### show oam ethernet connectivity-fault-management interfaces (trunk ports)

user@host> show oam ethernet connectivity-fault-management interfaces

| Interface                  | Link | Status   | Level | MEP Identifier | Neighbors |
|----------------------------|------|----------|-------|----------------|-----------|
| ge-4/0/1.0, vlan 100       | Up   | Active   | 5     | 100            | 0         |
| ge-10/3/10.4091, vlan 4091 | Down | Inactive | 4     | 400            | 0         |
| ge-4/0/0.0                 | Up   | Active   | 6     | 200            | 0         |

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/0.0

| Interface  | Link | Status | Level | MEP Identifier | Neighbors |
|------------|------|--------|-------|----------------|-----------|
| ge-4/0/0.0 | Up   | Active | 6     | 200            | 0         |

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vlan 100

| Interface            | Link | Status | Level | MEP Identifier | Neighbors |
|----------------------|------|--------|-------|----------------|-----------|
| ge-4/0/1.0, vlan 100 | Up   | Active | 5     | 100            | 0         |

user@host> **show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091 vlan 4091**

| Interface                  | Link | Status   | Level | MEP Identifier | Neighbors |
|----------------------------|------|----------|-------|----------------|-----------|
| ge-10/3/10.4091, vlan 4091 | Down | Inactive | 4     | 400            | 0         |

# show oam ethernet connectivity-fault-management linktrace path-database

## Syntax

```
show oam ethernet connectivity-fault-management linktrace path-database mac-address maintenance-association
ma-name maintenance-domain md-name
```

## Release Information

Command introduced in Junos OS Release 9.0.

## Description

On M320, MX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management maintenance linktrace database information.

## Options

**mac-address**—Display connectivity fault management path database information for the specified MAC address of the remote host.

**maintenance-association *ma-name***—Display connectivity fault management path database information for the specified maintenance association.

**maintenance-domain *md-name***—Display connectivity fault management path database information for the specified maintenance domain.

## Required Privilege Level

view

## List of Sample Output

[show oam ethernet connectivity-fault-management linktrace path-database on page 1894](#)  
[show oam ethernet connectivity-fault-management linktrace path-database \(Two traceroute Commands\) on page 1895](#)

## Output Fields

[Table 166 on page 1893](#) lists the output fields for the **show oam ethernet connectivity-fault-management linktrace path-database** command. Output fields are listed in the approximate order in which they appear.

Table 166: show oam ethernet connectivity-fault-management linktrace path-database Output Fields

| Field Name         | Field Description                                                           |
|--------------------|-----------------------------------------------------------------------------|
| Linktrace to       | MAC address of the 802.1ag node to which the linktrace message is targeted. |
| Interface          | Interface used by the local MEP to send the linktrace message (LTM).        |
| Maintenance Domain | Maintenance domain identifier specified in the traceroute command.          |

**Table 166: show oam ethernet connectivity-fault-management linktrace path-database Output Fields (continued)**

| Field Name                     | Field Description                                                                                                                                                                                                                                                             |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Maintenance Association</b> | Maintenance association identifier specified in the traceroute command.                                                                                                                                                                                                       |
| <b>Level</b>                   | Maintenance domain level configured for the maintenance domain.                                                                                                                                                                                                               |
| <b>Local Mep</b>               | MEP identifier of the local MEP originating the linktrace.                                                                                                                                                                                                                    |
| <b>Hop</b>                     | Sequential hop count of the linktrace path.                                                                                                                                                                                                                                   |
| <b>TTL</b>                     | Number of hops remaining in the linktrace message (LTM). 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 the next-hop MAC address derived from the next egress identifier in the Egress-ID TLV of the LTR PDU.                                                                                        |
| <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 responses (LTR), with a previously sent LTM. |

## Sample Output

**show oam ethernet connectivity-fault-management linktrace path-database**

**user@host> show oam ethernet connectivity-fault-management linktrace path-database  
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

```

**show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands)**

```

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:100002
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

Transaction Identifier:100003
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

```



## show oam ethernet connectivity-fault-management loss-statistics

### Syntax

```
show oam ethernet connectivity-fault-management loss-statistics
maintenance-domain md-name
maintenance-association ma-name
<count entry-count>
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

### Release Information

Command introduced in Junos OS Release 11.1.

### Description

On MX Series and ACX series routers with Ethernet interfaces, display ETH-LM statistics for on-demand mode only.

### Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**count *entry-count***—(Optional) Number of entries to display from the statistics table. The range of values is from 1 through 100. The default value is 100.

**local-mep *local-mep-id***—(Optional) Numeric identifier of the local MEP. The range of values is from 1 through 8191.

**remote-mep *remote-mep-id***—(Optional) Numeric identifier of the remote MEP. The range of values is from 1 through 8191.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[show oam ethernet connectivity-fault-management mep-statistics](#) | **1916**

### Output Fields

[Table 167 on page 1897](#) lists the output fields for the **show oam ethernet connectivity-fault-management loss-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 167: show oam ethernet connectivity-fault-management loss-statistics Output Fields

| Output Field Name                    | Field Description                                                                                                                                                              |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MEP identifier</b>                | Maintenance association end point (MEP) numeric identifier.                                                                                                                    |
| <b>MAC address</b>                   | Unicast MAC address configured for the MEP.                                                                                                                                    |
| <b>Remote MEP count</b>              | Number of remote MEPs (unless you specify the <b>remote-mep</b> option).                                                                                                       |
| <b>Remote MEP identifier</b>         | Numeric identifier of the remote MEP.                                                                                                                                          |
| <b>Remote MAC address</b>            | Unicast MAC address of the remote MEP.                                                                                                                                         |
| <b>Index</b>                         | Index number that corresponds to the ETH-LM entry in the CFM database.                                                                                                         |
| <b>Near-end frame loss</b>           | Count of frame loss associated with ingress data frames.                                                                                                                       |
| <b>Far-end frame loss</b>            | Count of frame loss associated with egress data frames.                                                                                                                        |
| <b>Near-end loss ratio</b>           | Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the ingress interface. |
| <b>Far-end loss ratio</b>            | Ratio, expressed as a percentage, of the number of service frames not delivered divided by the total number of service frames during time interval T at the egress interface.  |
| <b>Average near-end frame loss</b>   | Average frame loss measured in this session associated with ingress data frames.                                                                                               |
| <b>Average near-end loss ratio</b>   | Average frame loss ratio measured in this session associated with ingress data frames.                                                                                         |
| <b>Average far-end frame loss</b>    | Average frame loss measured in this session associated with egress data frames.                                                                                                |
| <b>Average far-end loss ratio</b>    | Average frame loss ratio measured in this session associated with egress data frames.                                                                                          |
| <b>Near-end best case loss</b>       | Lowest frame loss measured in this session associated with ingress data frames.                                                                                                |
| <b>Near-end best case loss ratio</b> | Lowest frame loss ratio measured in this session associated with ingress data frames.                                                                                          |

Table 167: show oam ethernet connectivity-fault-management loss-statistics Output Fields (continued)

| Output Field Name              | Field Description                                                                      |
|--------------------------------|----------------------------------------------------------------------------------------|
| Near-end worst case loss       | Highest frame loss measured in this session associated with ingress data frames.       |
| Near-end worst case loss ratio | Highest frame loss ratio measured in this session associated with ingress data frames. |
| Far-end best case frame loss   | Lowest frame loss measured in this session associated with egress data frames.         |
| Far-end best case loss ratio   | Lowest frame loss ratio measured in this session associated with egress data frames.   |
| Far-end worst case loss        | Highest frame loss measured in this session associated with egress data frames.        |
| Far-end worst case loss ratio  | Highest frame loss ratio measured in this session associated with egress data frames.  |

# show oam ethernet connectivity fault management loss statistics

user@host>. show oam ethernet connectivity fault management loss statistics maintenance-domain md maintenance-association ma

```

MEP identifier: 1, MAC address: 64:87:88:f9:7d:1b
Remote MEP count: 1

Remote MAC address: 64:87:88:6a:da:94

LM client session-id:4843
CIR Loss measurement statistics:
Index  Near-end      Far-end      Near-end      Far-end      Near-end
Far-end
      Frame loss  Total tx    Total rx      Frame loss  Total tx
Total rx
      (CIR)      (CIR)      (CIR)      (CIR)      (CIR)
(CIR)
      1          0          245          245          0          244
      244

```

|     |   |     |     |   |     |
|-----|---|-----|-----|---|-----|
| 2   | 0 | 488 | 488 | 0 | 489 |
| 489 |   |     |     |   |     |
| 3   | 0 | 732 | 732 | 0 | 733 |
| 733 |   |     |     |   |     |
| 4   | 0 | 977 | 977 | 0 | 976 |
| 976 |   |     |     |   |     |

EIR Loss measurement statistics:

| Index    | Near-end   | Far-end  | Near-end | Far-end    | Near-end |
|----------|------------|----------|----------|------------|----------|
| Far-end  |            |          |          |            |          |
|          | Frame loss | Total tx | Total rx | Frame loss | Total tx |
| Total rx |            |          |          |            |          |
|          | (EIR)      | (EIR)    | (EIR)    | (EIR)      | (EIR)    |
| (EIR)    |            |          |          |            |          |

|      |   |      |      |   |      |
|------|---|------|------|---|------|
| 1    | 0 | 272  | 272  | 0 | 273  |
| 273  |   |      |      |   |      |
| 2    | 0 | 546  | 546  | 0 | 545  |
| 545  |   |      |      |   |      |
| 3    | 0 | 820  | 820  | 0 | 819  |
| 819  |   |      |      |   |      |
| 4    | 0 | 1092 | 1092 | 0 | 1093 |
| 1093 |   |      |      |   |      |

Total far-end Tx (CIR) : 977

Total near-end Rx (CIR) : 977

Total near-end loss(CIR) : 0

Total near-end loss ratio(CIR) : 0.00000%

Total near-end Tx (CIR) : 976

Total far-end Rx (CIR) : 976

Total far-end loss(CIR) : 0

Total far-end loss ratio(CIR) : 0.00000%

Average near-end loss(CIR) : 0.00000

Average near-end loss ratio(CIR) : 0.00000%

Average far-end loss(CIR) : 0.00000

Average far-end loss ratio(CIR) : 0.00000%

Near-end best case loss(CIR) : 0

Near-end best case loss ratio(CIR) : 0.00000%

Near-end worst case loss(CIR) : 0

Near-end worst case loss ratio(CIR) : 0.00000%

Far-end best case loss(CIR) : 0

Far-end best case loss ratio(CIR) : 0.00000%

Far-end worst case loss(CIR) : 0

Far-end worst case loss ratio(CIR) : 0.00000%

Total far-end Tx (EIR) : 1092

Total near-end Rx (EIR) : 1092

Total near-end loss(EIR) : 0

```
Total near-end loss ratio(EIR)      : 0.00000%
Total near-end Tx (EIR)              : 1093
Total far-end Rx (EIR)               : 1093
Total far-end loss(EIR)              : 0
Total far-end loss ratio(EIR)        : 0.00000%
Average near-end loss(EIR)           : 0.00000
Average near-end loss ratio(EIR)     : 0.00000%
Average far-end loss(EIR)            : 0.00000
Average far-end loss ratio(EIR)      : 0.00000%
Near-end best case loss(EIR)         : 0
Near-end best case loss ratio(EIR)   : 0.00000%
Near-end worst case loss(EIR)        : 0
Near-end worst case loss ratio(EIR)  : 0.00000%
Far-end best case loss(EIR)          : 0
Far-end best case loss ratio(EIR)    : 0.00000%
Far-end worst case loss(EIR)         : 0
Far-end worst case loss ratio(EIR)   : 0.00000%
```

## show oam ethernet connectivity-fault-management mep-database

### Syntax

```
show oam ethernet connectivity-fault-management mep-database
maintenance-domain domain-name
maintenance-association ma-name
<local-mep local-mep-id>
<remote-mep remote-mep-id>
extensive
```

### Release Information

Command introduced in Junos OS Release 8.4.

Support for ITU-T Y.1731 frame delay measurement added in Junos OS Release 9.5.

Support for ITU-T Y.1731 synthetic frame loss measurement added in Junos OS Release 13.2 for MX Series routers.

Support for ITU-T Y.1731 ethernet expected defect function (ETH-ED) added in Junos OS Release 19.1 for MX Series routers

**NOTE:** Ethernet expected defect function (ETH-ED) will be supported on all MX series routers only when **enhanced-ip** mode is enabled. ETH-ED will not be supported if CFM is running in centralized mode.

### Description

On M7i and M10i routers with Enhanced CFEB (CFEB-E), and on M320, M120, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

In addition, on M120, M320, and MX series routers, also display port status TLV, interface status TLV, and action profile information.

In addition, for Ethernet interfaces on MX Series routers, also display any ITU-T Y.1731 frame delay measurement (ETH-DM) frame counts.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

### Options

**maintenance-association *ma-name***—Name of the maintenance association.

**maintenance-domain *domain-name***—Name of the maintenance domain.

**local-mep-id**—(Optional) Numeric identifier of local MEP.

**remote-mep-id**—(Optional) Numeric identifier of the remote MEP.

**Required Privilege Level**

view

RELATED DOCUMENTATION

|                                                                                         |
|-----------------------------------------------------------------------------------------|
| <a href="#">clear oam ethernet connectivity-fault-management statistics   1444</a>      |
| <a href="#">show oam ethernet connectivity-fault-management delay-statistics   1869</a> |
| <a href="#">show oam ethernet connectivity-fault-management interfaces   1879</a>       |
| <a href="#">show oam ethernet connectivity-fault-management mep-statistics   1916</a>   |

**List of Sample Output**

- [show oam ethernet connectivity-fault-management mep-database on page 1908](#)
- [show oam ethernet connectivity-fault-management mep-database \(One-Way ETH-DM\) on page 1909](#)
- [show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 1909](#)
- [show oam ethernet connectivity-fault-management mep-database remote-mep \(Action Profile Event\) on page 1910](#)
- [show oam ethernet connectivity-fault-management mep-database \(Connection Protection TLV Configured\) on page 1910](#)
- [show oam ethernet connectivity-fault-management mep-database on page 1912](#)
- [show oam ethernet connectivity-fault-management mep-database \(enhanced continuity measurement\) on page 1913](#)
- [show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 1914](#)

**Output Fields**

Table 168 on page 1902 lists the output fields for the **show oam ethernet connectivity-fault-management mep-database** command. Output fields are listed in the approximate order in which they appear.

Table 168: show oam ethernet connectivity-fault-management mep-database Output Fields

| Field Name                  | Field Description                          |
|-----------------------------|--------------------------------------------|
| Maintenance domain name     | Maintenance domain name.                   |
| Format (Maintenance domain) | Maintenance domain name format configured. |
| Level                       | Maintenance domain level configured.       |

Table 168: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

| Field Name                              | Field Description                                                                                                                                                                                                                                                                                |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Maintenance association name</b>     | Maintenance association name.                                                                                                                                                                                                                                                                    |
| <b>Format (Maintenance association)</b> | Maintenance association name format configured.                                                                                                                                                                                                                                                  |
| <b>Continuity-check status</b>          | Continuity-check status.                                                                                                                                                                                                                                                                         |
| <b>Interval</b>                         | Continuity-check message interval.                                                                                                                                                                                                                                                               |
| <b>Loss-threshold</b>                   | Lost continuity-check message threshold.                                                                                                                                                                                                                                                         |
| <b>Connection Protection TLV</b>        | Status of the connection protection TLV, if configured on the MEP interface: <b>no</b> , <b>yes</b><br><br>If <b>yes</b> , then the transmitted connection protection TLV is decoded and the following three fields are displayed: <b>Prefer me</b> , <b>Protection in use</b> , <b>FRR Flag</b> |
| <b>Prefer me</b>                        | If set to <b>yes</b> , the path through which CCM was transmitted is preferred (unless the path fails). It is used for signaling a manual-switch command to remote side.<br><br>Its value can be <b>yes</b> or <b>no</b> .                                                                       |
| <b>Protection in use</b>                | Used for protection decision coordination. Its value is set to <b>yes</b> if the endpoint transmitting the CCM is currently transmitting the user traffic to protection path.<br><br>Its value can be <b>yes</b> or <b>no</b> .                                                                  |
| <b>FRR Flag</b>                         | LSR/LER forwarding the CCM Frame into a bypass tunnel is set.<br><br>Its value can be <b>yes</b> or <b>no</b> .                                                                                                                                                                                  |
| <b>MEP identifier</b>                   | Maintenance association end point (MEP) identifier.                                                                                                                                                                                                                                              |
| <b>Direction</b>                        | MEP direction configured.                                                                                                                                                                                                                                                                        |
| <b>MAC address</b>                      | MAC address configured for the MEP.                                                                                                                                                                                                                                                              |
| <b>Auto-discovery</b>                   | Whether automatic discovery is enabled or disabled.                                                                                                                                                                                                                                              |
| <b>Priority</b>                         | Priority used for CCMs and linktrace messages transmitted by the MEP.                                                                                                                                                                                                                            |
| <b>Interface name</b>                   | Interface identifier.                                                                                                                                                                                                                                                                            |



Table 168: show oam ethernet connectivity-fault-management mep-database Output Fields (continued)

| Field Name                           | Field Description                                                                                                                                                                                                 |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface status</b>              | Local interface status.                                                                                                                                                                                           |
| <b>Link status</b>                   | Local link status.                                                                                                                                                                                                |
| <b>Remote MEP not receiving CCM</b>  | Whether the remote MEP is not receiving CCMs.                                                                                                                                                                     |
| <b>Erroneous CCM received</b>        | Whether erroneous CCMs have been received.                                                                                                                                                                        |
| <b>Cross-connect CCM received</b>    | Whether cross-connect CCMs have been received.                                                                                                                                                                    |
| <b>RDI sent by some MEP</b>          | Whether the remote defect indication (RDI) bit is set in messages that have been received.<br>The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs. |
| <b>CCMs sent</b>                     | Number of CCMs transmitted.                                                                                                                                                                                       |
| <b>CCMs received out of sequence</b> | Number of CCMs received out of sequence.                                                                                                                                                                          |
| <b>LBMs sent</b>                     | Number of loopback messages (LBMs) sent.                                                                                                                                                                          |
| <b>Valid in-order LBRs received</b>  | Number of loopback response messages (LBRs) received that were valid messages and in sequence.                                                                                                                    |
| <b>1DMs sent</b>                     | If the MEP is an initiator for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.<br><br>For all other cases, this field displays 0.            |
| <b>Valid 1DMs received</b>           | If the MEP is a receiver for a one-way ETH-DM session: Number of valid 1DM frames received.<br><br>For all other cases, this field displays 0.                                                                    |
| <b>Invalid 1DMs received</b>         | If the MEP is a receiver for a one-way ETH-DM session: Number of invalid 1DM frames received.<br><br>For all other cases, this field displays 0.                                                                  |
| <b>Out of sync 1DMs received</b>     | If the MEP is a receiver for a one-way ETH-DM session: Number of out-of-sync one-way delay measurement request packets received.                                                                                  |

Table 168: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

| Field Name                                 | Field Description                                                                                                                                                                                      |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DMMs sent</b>                           | If the MEP is an initiator for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.<br><br>For all other cases, this field displays 0. |
| <b>Valid DMMs received</b>                 | If the MEP is an initiator for a two-way ETH-DM session: Number of valid two-way delay measurement packets received.                                                                                   |
| <b>Invalid DMMs received</b>               | If the MEP is an initiator for a two-way ETH-DM session: Number of invalid two-way delay measurement packets received.                                                                                 |
| <b>DMRs sent</b>                           | If the MEP is a responder for a ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent.<br><br>For all other cases, this field displays 0.                                                |
| <b>Valid DMRs received</b>                 | If the MEP is an initiator for a two-way ETH-DM session: Number of valid DMRs received.<br><br>For all other cases, this field displays 0.                                                             |
| <b>Invalid DMRs received</b>               | If the MEP is an initiator for a two-way ETH-DM session: Number of invalid DMRs received.<br><br>For all other cases, this field displays 0.                                                           |
| <b>Valid out-of-order LBRs received</b>    | Number of LBRs received that were valid messages and not in sequence.                                                                                                                                  |
| <b>LBRs received with corrupted data</b>   | Number of LBRs received that were corrupted.                                                                                                                                                           |
| <b>LBRs sent</b>                           | Number of LBRs transmitted.                                                                                                                                                                            |
| <b>LTM sent</b>                            | Linktrace messages (LTMs) transmitted.                                                                                                                                                                 |
| <b>LTM received</b>                        | Linktrace messages received.                                                                                                                                                                           |
| <b>LTR sent</b>                            | Linktrace responses (LTRs) transmitted.                                                                                                                                                                |
| <b>LTR received</b>                        | Linktrace responses received.                                                                                                                                                                          |
| <b>Sequence number of next LTM request</b> | Sequence number of the next linktrace message request to be transmitted.                                                                                                                               |

Table 168: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

| Field Name                   | Field Description                                                                                                                                                                                                    |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>LMM sent</b>              | If the interface is attached to an initiator MEP for a ETH-LM session: Number of loss measurement message (LMM) PDU frames sent to the peer MEP in this session.                                                     |
| <b>Valid LMM received</b>    | If the interface is attached to an initiator MEP for a ETH-LM session: Number of valid loss measurement request packets received.                                                                                    |
| <b>Invalid LMM received</b>  | If the interface is attached to an initiator MEP for a ETH LM session: Number of invalid loss measurement request packets received.                                                                                  |
| <b>LMR sent</b>              | If the interface is attached to a responder MEP for a ETH-LM session: Number of loss measurement reply (LMR) frames sent.                                                                                            |
| <b>Valid LMR received</b>    | If the interface is attached to an initiator MEP for a ETH LM session: Number of valid LMR frames received.                                                                                                          |
| <b>Invalid LMR received</b>  | If the interface is attached to an initiator MEP for a ETH-LM session: Number of invalid LMR frames received.                                                                                                        |
| <b>SLM sent</b>              | If the interface is attached to an initiator MEP for a ETH-SLM session: Number of synthetic loss measurement (SLM) request packets transmitted from the source MEP to the remote or destination MEP in this session. |
| <b>Valid SLM received</b>    | If the interface is attached to a responder MEP for a ETH-SLM session: Number of valid SLM PDUs transmitted from the source MEP to the remote or destination MEP.                                                    |
| <b>Invalid SLM received</b>  | If the interface is attached to a responder MEP for a ETH-SLM session: Number of invalid SLM PDUs transmitted from the source MEP to the remote or destination MEP.                                                  |
| <b>SLR sent</b>              | If the interface is attached to a responder MEP for a ETH-SLM session: Number detail extensive of synthetic loss reply (SLR) frames sent.                                                                            |
| <b>Valid SLR received</b>    | If the interface is attached to an initiator MEP for a ETH-SLM session: Number of valid SLR PDUs that the source MEP received from the remote or destination MEP.                                                    |
| <b>Invalid SLR received</b>  | If the interface is attached to an initiator MEP for a ETH-SLM session: Number of invalid SLR PDUs that the source MEP received from the remote or destination MEP.                                                  |
| <b>Remote MEP identifier</b> | MEP identifier of the remote MEP.                                                                                                                                                                                    |
| <b>State (remote MEP)</b>    | State of the remote MEP: <b>idle</b> , <b>start</b> , <b>ok</b> , or <b>failed</b> .                                                                                                                                 |

Table 168: show oam ethernet connectivity-fault-management mep-database Output Fields (continued)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MAC address              | MAC address of the remote MEP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Type                     | Whether the remote MEP MAC address was learned using automatic discovery or configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Interface                | Interface of the remote MEP. A seven-digit number is appended if CFM is configured to run on a routing instance of type VPLS.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Last flapped             | Date, time, and how long ago the remote MEP 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> .                                                                                                                                                                                                                                                                              |
| Remote defect indication | Whether the remote defect indication (RDI) bit is set in messages that have been received or transmitted.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Port status TLV          | <ul style="list-style-type: none"> <li>In the Maintenance domain section, displays the last transmitted port status TLV value.</li> <li>In the Remote MEP section, displays the last value of port status TLV received from the remote MEP.</li> </ul> <p>In the Action profile section, displays, the last occurred event <b>port-status-tlv blocked</b> event. This event occurred due to the reception of <b>blocked</b> value in the port status TLV from remote MEP.</p>                                                                      |
| Interface status TLV     | <ul style="list-style-type: none"> <li>In the Maintenance domain section, displays the last transmitted interface status TLV value.</li> <li>In the Remote MEP section, displays the last value of interface status TLV received from the remote MEP.</li> </ul> <p>In the Action profile section, if displays, the last occurred event interface-status-tlv event (either <b>lower-layer-down</b> or <b>down</b>). This event occurred due to the reception of either lower or <b>down</b> value in the interface status TLV from remote MEP.</p> |
| Action profile           | Name of the action profile occurrence associated with a remote MEP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Last event               | When an action profile occurs, displays the last event that triggered it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Last event cleared       | When all the configured and occurred events (under action profile) are cleared, then the action taken gets reverted (such as down interface is made up) and the corresponding time is noted and displayed.                                                                                                                                                                                                                                                                                                                                         |
| Action                   | Action taken and the corresponding time of the action occurrence.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

**show oam ethernet connectivity-fault-management mep-database**

**user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain  
vpls-vlan2000 maintenance-association vpls-vlan200**

```
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no Interface name: ge-0/0/1.0, Interface status:
Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMs sent                                  : 1476
  CCMs received out of sequence              : 0
  LBMs sent                                  : 85
  Valid in-order LBRs received               : 78
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 1
  LTMs received                             : 0
  LTRs sent                                  : 0
  LTRs received                             : 1
  Sequence number of next LTM request        : 1
  lDMs sent                                  : 0
  Valid lDMs received                       : 0
  Invalid lDMs received                     : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                       : 0
  Invalid DMRs received                     : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    100      00:19:e2:b2:81:4b  ok    vt-0/1/10.1049088
```

# show oam ethernet connectivity-fault-management mep-database (One-Way ETH-DM)

user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md6 maintenance-domain ma6

```
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMs sent                                  : 1590
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                              : 0
  LTRs sent                                  : 0
  LTRs received                              : 0
  Sequence number of next LTM request        : 0
  lDMs sent                                  : 10
  Valid lDMs received                       : 0
  Invalid lDMs received                     : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                      : 0
  Invalid DMRs received                    : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    201      00:90:69:0a:43:94  ok    ge-0/2/5.0
```

# show oam ethernet connectivity-fault-management mep-database local-mep remote-mep

user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200 local-mep 200 remote-mep 100

```

Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

Remote MEP identifier: 100, State: ok
MAC address: 00:19:e2:b2:81:4b, Type: Learned
Interface: vt-0/1/10.1049088
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none

```

**show oam ethernet connectivity-fault-  
management mep-database remote-mep  
(Action Profile Event)**

```

user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md5 maintenance-association ma5 remote-mep 200

```

```

Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

Remote MEP identifier: 200, State: ok
MAC address: 00:05:85:73:96:1f, Type: Configured
Interface: ge-1/0/8.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: lower-layer-down
Action profile: juniper
  Last event: Interface-status-tlv lower-layer-down
  Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)

```

**show oam ethernet connectivity-fault-  
management mep-database**

**(Connection Protection TLV Configured)**

```
user@host>show oam ethernet connectivity-fault-management mep-database maintenance-domain
md5 maintenance-association ma5
```

If connection-protection is not enabled on down MEPs, but connection-protection TLV is used, MX always sets the protection-in-use flag in connection-protection tlv, while CCMs are sent out. During reversion, this is an indicator to the receiver that protect-path is in use, otherwise the peer (receiver) assumes working is active and reversion does not work as expected. Setting this bit does not affect protection-switching/traffic-loss.

```
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:19:e2:b1:14:30
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: yes
  Prefer me: no, Protection in use: no, FRR Flag: no
Interface name: xe-6/2/0.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : no
  Cross-connect CCM received                   : no
  RDI sent by some MEP                        : no
  Some remote MEP's MAC in error state         : no
Statistics:
  CCMs sent                                   : 251
  CCMs received out of sequence                : 0
  LBMs sent                                   : 0
  Valid in-order LBRs received                 : 0
  Valid out-of-order LBRs received             : 0
  LBRs received with corrupted data            : 0
  LBRs sent                                   : 0
  LTMs sent                                   : 0
  LTMs received                               : 0
  LTRs sent                                   : 0
  LTRs received                               : 0
  Sequence number of next LTM request          : 0
  1DMs sent                                   : 0
  Valid 1DMs received                         : 0
  Invalid 1DMs received                       : 0
  Out of sync 1DMs received                   : 0
  DMMs sent                                   : 0
  Valid DMMs received                         : 0
```



```

Invalid DMMS received           : 0
DMRs sent                       : 0
Valid DMRs received             : 0
Invalid DMRs received           : 0
LMMs sent                       : 0
Valid LMMs received             : 0
Invalid LMMs received           : 0
LMRs sent                       : 0
Valid LMRs received             : 0
Invalid LMRs received           : 0
Remote MEP count: 1
Identifier    MAC address      State    Interface
      2      00:90:69:7f:e4:30

```

### **show oam ethernet connectivity-fault-management mep-database**

**user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain md5 maintenance-association ma5**

```

Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:14:f6:b6:01:fe
Auto-discovery: enabled, Priority: 0
Interface name: ge-1/0/0.0, Interface status: Active, Link status: Up

Defects:
Remote MEP not receiving CCM      : no
Erroneous CCM received            : no
Cross-connect CCM received        : no
RDI sent by some MEP              : no

Statistics:
CCMs sent                        : 328703
CCMs received out of sequence     : 0
LBMs sent                        : 85
Valid in-order LBRs received      : 78
Valid out-of-order LBRs received  : 0
LBRs received with corrupted data : 0
LBRs sent                        : 0
LTMs sent                        : 0
LTMs received                     : 0
LTRs sent                        : 0
LTRs received                     : 0

```

```

Sequence number of next LTM request : 0
1DMs sent                          : 10
Valid 1DMs received                 : 10
Invalid 1DMs received                : 0
DMMs sent                          : 20
DMRs sent                          : 0
Valid DMRs received                 : 10
Invalid DMRs received                : 0
LMM sent                          : 10
Valid LMM received                  : 20
Invalid LMM received                 : 0
LMR sent                          : 20
Valid LMR received                  : 10
Invalid LMR received                : 0
SLM sent                          : 10
Valid SLM received                  : 20
Invalid SLM received                 : 0
SLR sent                          : 20
Valid SLR received                  : 10
Invalid SLR received                 : 0

Remote MEP count                    : 1

Identifier      MAC address      State      Interface
    2          00:12:1e:fb:ea:7d   ok        ge-1/0/0.0

```

**show oam ethernet connectivity-fault-  
management mep-database (enhanced continuity measurement)**

**user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain  
md5 maintenance-association ma5 local-mep 2001 remote-mep 1001**

```

Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address   : 00:19:e2:b0:74:00, Type: Learned
Interface     : ge-2/0/0.0
Last flapped  : Never

```

```
+ Continuity : 91%, Admin-enable duration: 2100sec, Oper-down duration: 100sec
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
```

### **show oam ethernet connectivity-fault-management mep-database local-mep remote-mep**

When the timer is not running:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md6 maintenance-association ma6 local-mep 201 remote-mep 101 extensive
```

```
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 201, Direction: up, MAC address: 00:24:dc:48:46:cd
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no
Chassis ID: 00:24:dc:48:47:f2
Management Address: 10.216.34.30
Sendid TLV: no
Interface name: ge-5/1/9.0, Interface status: Active, Link status: Up

Remote MEP identifier: 101, State: ok
MAC address: 00:22:83:db:4c:5b, Type: Configured
Interface: ge-5/3/0.0
Last flapped: 2019-02-05 00:45:57 PST (00:54:51 ago)
Continuity: 99%, Admin-enable duration: 3367sec, Oper-down duration: 43sec
Effective loss threshold: 3 frames
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
Remote Management Address: 0.0.0.0
Sendid TLV: no
Sender ID Remote Chassis Subtype: 0
EDM time remaining: N/A
```

When the timer is running:

```
user@host> show oam ethernet connectivity-fault-management mep-database maintenance-domain
md6 maintenance-association ma6 local-mep 201 remote-mep 101 extensive
```

```
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 201, Direction: up, MAC address: 00:24:dc:48:46:cd
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Connection Protection TLV: no
Chassis ID: 00:24:dc:48:47:f2
Management Address: 10.216.34.30
Sendid TLV: no
Interface name: ge-5/1/9.0, Interface status: Active, Link status: Up

Remote MEP identifier: 101, State: ok
MAC address: 00:22:83:db:4c:5b, Type: Configured
Interface: ge-5/3/0.0
Last flapped: 2019-02-05 00:45:57 PST (00:56:16 ago)
Continuity: 99%, Admin-enable duration: 3451sec, Oper-down duration: 43sec
Effective loss threshold: 3 frames
Remote defect indication: false
Port status TLV: none
Interface status TLV: none
Remote Management Address: 0.0.0.0
Sendid TLV: no
Sender ID Remote Chassis Subtype: 0
EDM time remaining: 894 secs
```

## show oam ethernet connectivity-fault-management mep-statistics

### Syntax

```
show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md-name
maintenance-association ma-name
<mep mep-id>
<remote-mep remote-mep-id>
<count entry-count>
```

### Release Information

Command introduced in Junos OS Release 9.5.

Command introduced in Junos OS Release 11.4 for EX Series switches.

Support for ITU-T Y.1731 Ethernet synthetic frame loss measurement (ETH-SLM) added in Junos OS Release 13.2 for MX Series routers.

### Description

On MX Series and ACX Series routers and EX Series switches with Ethernet interfaces, display ETH-DM statistics and ETH-DM frame counts.

For Ethernet interfaces on MX Series routers, display any ITU-T Y.1731 synthetic frame loss measurement (ETH-SLM) statistics and frame counts.

### Options

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**mep *mep-id***—(Optional) Numeric identifier of the local MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.

**remote-mep *remote-mep-id***—(Optional) Numeric identifier of the remote MEP. The range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.

**count *entry-count***—(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[clear oam ethernet connectivity-fault-management statistics](#) | [1444](#)

---

[show oam ethernet connectivity-fault-management delay-statistics | 1869](#)


---

[show oam ethernet connectivity-fault-management interfaces | 1879](#)


---

[show oam ethernet connectivity-fault-management mep-database | 1901](#)


---

### List of Sample Output

[show oam ethernet connectivity-fault-management mep-statistics \(CIR counters only\) on page 1919](#)
[show oam ethernet connectivity-fault-management mep-statistics \(CIR and EIR counters enabled\) on page 1921](#)
[show oam ethernet connectivity-fault-management mep-statistics remote-mep \(CIR counters only\) on page 1923](#)
[show oam ethernet connectivity-fault-management mep-statistics remote-mep \(CIR and EIR counters enabled\) on page 1925](#)
[show oam ethernet connectivity-fault-management mep-statistics on page 1927](#)
[show oam ethernet connectivity-fault-management mep-statistics remote-mep on page 1928](#)
[show oam ethernet connectivity-fault-management mep-statistics local-mep remote-mep on page 1930](#)

### Output Fields

[Table 169 on page 1917](#) lists the output fields for the **show oam ethernet connectivity-fault-management mep-statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 169: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields**

| Output Field Name             | Field Description                                                                                          |
|-------------------------------|------------------------------------------------------------------------------------------------------------|
| MEP identifier                | Maintenance association end point (MEP) numeric identifier.                                                |
| MAC address                   | Unicast MAC address configured for the MEP.                                                                |
| Remote MEP count              | Number of remote MEPs (unless you specify the <b>remote-mep</b> option).                                   |
| CCMs sent                     | Number of continuity check messages (CCMs) sent.                                                           |
| CCMs received                 | Number of continuity check messages (CCMs) received for a specific remote MEP and maintenance association. |
| CCMs received out of sequence | Number of continuity check messages (CCMs) received that were not in sequence.                             |
| Remote MEP identifier         | Numeric identifier of the remote MEP.                                                                      |
| Remote MAC address            | Unicast MAC address of the remote MEP.                                                                     |
| Index                         | Index number that corresponds to the ETH-DM entry in the CFM database.                                     |

**Table 169: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields (continued)**

| Output Field Name                      | Field Description                                                                                                                                                                                                                                                                                                        |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>One-way delay (usec)</b>            | <p>For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP.</p> <p>For a detailed description of one-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>  |
| <b>Two-way delay (usec)</b>            | <p>For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP.</p> <p>For a detailed description of two-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p> |
| <b>Average one-way delay</b>           | Average one-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                |
| <b>Average one-way delay variation</b> | Average one-way “frame jitter” for the statistics displayed.                                                                                                                                                                                                                                                             |
| <b>Best-case one-way delay</b>         | Lowest one-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                 |
| <b>Worst-case one-way delay</b>        | Highest one-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                |
| <b>Average two-way delay</b>           | Average two-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                |
| <b>Average two-way delay variation</b> | Average two-way “frame jitter” for the statistics displayed.                                                                                                                                                                                                                                                             |
| <b>Best-case two-way delay</b>         | Lowest two-way frame delay for the statistics displayed.                                                                                                                                                                                                                                                                 |
| <b>Worst-case two-way delay</b>        | Highest two-way frame delay calculated in this session.                                                                                                                                                                                                                                                                  |
| <b>SLM packets sent</b>                | Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.                                                                                                                                                                                          |
| <b>SLM packets received</b>            | Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.                                                                                                                                                                                    |
| <b>SLR packets sent</b>                | Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.                                                                                                                                                                                        |
| <b>SLR packets received</b>            | Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.                                                                                                                                                                                  |

**Table 169: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields (continued)**

| Output Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Local TXFC1 value</b>                 | Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.   |
| <b>Local RXFC1 value</b>                 | Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet. |
| <b>Last Received SLR frame TXFCf(tc)</b> | Value of the local counter TxFCI at the time of SLM frame transmission.                                                                                                                                                                                                                                                           |
| <b>Last Received SLR frame TXFCb(t)</b>  | Value of the local counter RxFCI at the time of SLR frame transmission.                                                                                                                                                                                                                                                           |
| <b>Frame loss (near-end)</b>             | Count of frame loss associated with ingress data frames.                                                                                                                                                                                                                                                                          |
| <b>Frame loss (far-end)</b>              | Count of frame loss associated with egress data frames.                                                                                                                                                                                                                                                                           |

## Sample Output

**show oam ethernet connectivity-fault-management mep-statistics (CIR counters only)**

**user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3**

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count           : 1
CCMs sent                  : 6550
CCMs received out of sequence : 0
LBMs sent                  : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent                  : 0
LTMs sent                  : 0
```



```

LTMs received           : 0
LTRs sent               : 0
LTRs received           : 0
Sequence number of next LTM request : 0
1DMs sent               : 5
Valid 1DMs received     : 0
Invalid 1DMs received   : 0
DMMs sent               : 5
DMRs sent               : 0
Valid DMRs received     : 5
Invalid DMRs received   : 0
LMM sent                : 5
Valid LMM received      : 5
Invalid LMM received    : 0
LMR sent                : 0
Valid LMR received      : 5
Invalid LMR received    : 0
Remote MEP identifier    : 101
Remote MAC address       : 00:05:85:73:39:4a

```

Delay measurement statistics:

| Index | One-way delay<br>(usec) | Two-way delay<br>(usec) |
|-------|-------------------------|-------------------------|
| 1     | 259                     | 519                     |
| 2     | 273                     | 550                     |
| 3     | 287                     | 571                     |
| 4     | 299                     | 610                     |
| 5     | 313                     | 650                     |

```

Average one-way delay      : 286 usec
Average one-way delay variation : 62 usec
Best case one-way delay    : 259 usec
Average two-way delay      : 580 usec
Average two-way delay variation : 26 usec
Best case two-way delay    : 519 usec
Worst case two-way delay   : 650 usec

```

Loss measurement statistics:

| Index | Near-end<br>Frame loss<br>(CIR) | Far-end<br>Frame loss<br>(CIR) | Near-end<br>Frame loss<br>(EIR) | Far-end<br>Frame loss<br>(EIR) |
|-------|---------------------------------|--------------------------------|---------------------------------|--------------------------------|
| 1     | 9                               | 9                              |                                 |                                |
| 2     | 3                               | 5                              |                                 |                                |
| 3     | 7                               | 5                              |                                 |                                |

|   |   |   |
|---|---|---|
| 4 | 9 | 6 |
| 5 | 3 | 6 |

```

Average near-end loss (CIR)           : 6.2
Average near-end loss ratio (CIR)      : 6.2%
Average far-end loss (CIR)            : 6.2
Average far-end loss ratio (CIR)       : 6.2%
Near-end best case loss (CIR)         : 3
Near-end best case loss ratio (CIR)    : 3%
Near-end worst case loss (CIR)        : 9
Near-end worst case loss ratio (CIR)   : 9%
Far-end best case loss (CIR)          : 5
Far-end best case loss ratio (CIR)     : 5%
Far-end worst case loss (CIR)         : 9
Far-end worst case loss ratio (CIR)    : 9%

```

#### **show oam ethernet connectivity-fault-management mep-statistics (CIR and EIR counters enabled)**

**user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3**

```

MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count           : 1
CCMs sent                  : 6550
CCMs received out of sequence : 0
LBMs sent                  : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent                  : 0
LTMs sent                  : 0
LTMs received              : 0
LTRs sent                  : 0
LTRs received              : 0
Sequence number of next LTM request : 0
1DMs sent                  : 5
Valid 1DMs received        : 0
Invalid 1DMs received       : 0
DMMs sent                  : 5
DMRs sent                  : 0
Valid DMRs received        : 5
Invalid DMRs received       : 0
LMM sent                   : 5

```

```

alid LMM received           : 5
Invalid LMM received        : 0
LMR sent                    : 0
Valid LMR received          : 5
Invalid LMR received        : 0
Remote MEP identifier       : 101
Remote MAC address          : 00:05:85:73:39:4a

```

Delay measurement statistics:

| Index | One-way delay<br>(usec) | Two-way delay<br>(usec) |
|-------|-------------------------|-------------------------|
| 1     | 259                     | 519                     |
| 2     | 273                     | 550                     |
| 3     | 287                     | 571                     |
| 4     | 299                     | 610                     |
| 5     | 313                     | 650                     |

```

Average one-way delay           : 286 usec
Average one-way delay variation : 62 usec
Best case one-way delay         : 259 usec
Average two-way delay           : 580 usec
Average two-way delay variation : 26 usec
Best case two-way delay         : 519 usec
Worst case two-way delay        : 650 usec

```

Loss measurement statistics:

| Index | Near-end<br>Frame loss<br>(CIR) | Far-end<br>Frame loss<br>(CIR) | Near-end<br>Frame loss<br>(EIR) | Far-end<br>Frame loss<br>(EIR) |
|-------|---------------------------------|--------------------------------|---------------------------------|--------------------------------|
| 1     | 9                               | 9                              | 2                               | 4                              |
| 2     | 3                               | 5                              | 4                               | 6                              |
| 3     | 7                               | 5                              | 0                               | 2                              |
| 4     | 9                               | 6                              | 8                               | 2                              |
| 5     | 3                               | 6                              | 6                               | 4                              |

```

Average near-end loss (CIR)           : 6.2
Average near-end loss ratio (CIR)     : 6.2%
Average far-end loss (CIR)            : 6.2
Average far-end loss ratio (CIR)      : 6.2%
Near-end best case loss (CIR)         : 3
Near-end best case loss ratio (CIR)   : 3%
Near-end worst case loss (CIR)        : 9
Near-end worst case loss ratio (CIR)  : 9%
Far-end best case loss (CIR)          : 5

```

```

Far-end best case loss ratio (CIR)      : 5%
Far-end worst case loss (CIR)          : 9
Far-end worst case loss ratio (CIR)    : 9%
Average near-end loss (EIR)            : 4
Average near-end loss ratio (EIR)      : 4%
Average far-end loss (EIR)             : 3.4
Average far-end loss ratio (EIR)       : 3.4%
Near-end best case loss (EIR)          : 0
Near-end best case loss ratio (EIR)    : 0%
Near-end worst case loss (EIR)         : 8
Near-end worst case loss ratio (EIR)   : 8%
Far-end best case loss (EIR)           : 2
Far-end best case loss ratio (EIR)     : 2%
Far-end worst case loss (EIR)          : 6
Far-end worst case loss ratio (EIR)    : 6%

```

**show oam ethernet connectivity-fault-  
management mep-statistics remote-mep (CIR counters only)**

user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain  
md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3 remote-mep 101**

```

MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                      : 7762
CCMs received out of sequence  : 0
LBMs sent                      : 0
Valid in-order LBRs received   : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent                     : 0
LTMs sent                     : 0
LTMs received                  : 0
LTRs sent                     : 0
LTRs received                  : 0
Sequence number of next LTM request : 0
LDMs sent                     : 5
Valid LDMs received            : 0
Invalid LDMs received          : 0
DMMs sent                     : 5
DMRs sent                     : 0
Valid DMRs received            : 5
Invalid DMRs received          : 0
LMM sent                      : 5
Valid LMM received             : 5

```

```

Invalid LMM received      : 0
LMR sent                  : 0
Valid LMR received        : 5
Invalid LMR received      : 0
Remote MEP identifier      : 101
Remote MAC address        : 00:05:85:73:39:4a

```

Delay measurement statistics:

| Index | One-way delay<br>(usec) | Two-way delay<br>(usec) |
|-------|-------------------------|-------------------------|
| 1     | 259                     | 519                     |
| 2     | 273                     | 550                     |
| 3     | 287                     | 571                     |
| 4     | 299                     | 610                     |
| 5     | 313                     | 650                     |

```

Average one-way delay      : 286 usec
Average one-way delay variation : 62 usec
Best case one-way delay    : 259 usec
Average two-way delay      : 580 usec
Average two-way delay variation : 26 usec
Best case two-way delay    : 519 usec
Worst case two-way delay   : 650 usec

```

Loss measurement statistics:

| Index | Near-end<br>Frame loss<br>(CIR) | Far-end<br>Frame loss<br>(CIR) | Near-end<br>Frame loss<br>(EIR) | Far-end<br>Frame loss<br>(EIR) |
|-------|---------------------------------|--------------------------------|---------------------------------|--------------------------------|
| 1     | 9                               | 9                              |                                 |                                |
| 2     | 3                               | 5                              |                                 |                                |
| 3     | 7                               | 5                              |                                 |                                |
| 4     | 9                               | 6                              |                                 |                                |
| 5     | 3                               | 6                              |                                 |                                |

```

Average near-end loss (CIR)      : 6.2
Average near-end loss ratio (CIR) : 6.2%
Average far-end loss (CIR)      : 6.2
Average far-end loss ratio (CIR) : 6.2%
Near-end best case loss (CIR)    : 3
Near-end best case loss ratio (CIR) : 3%
Near-end worst case loss (CIR)   : 9
Near-end worst case loss ratio (CIR) : 9%
Far-end best case loss (CIR)     : 5
Far-end best case loss ratio (CIR) : 5%
Far-end worst case loss (CIR)    : 9

```

```

Far-end worst case loss ratio (CIR)      : 9%
Average near-end loss (EIR)              : 4
Average near-end loss ratio (EIR)        : 4%
Average far-end loss (EIR)               : 3.4
Average far-end loss ratio (EIR)         : 3.4%
Near-end best case loss (EIR)            : 0
Near-end best case loss ratio (EIR)      : 0%
Near-end worst case loss (EIR)           : 8
Near-end worst case loss ratio (EIR)     : 8%
Far-end best case loss (EIR)             : 2
Far-end best case loss ratio (EIR)       : 2%
Far-end worst case loss (EIR)            : 6
Far-end worst case loss ratio (EIR)      : 6%

```

**show oam ethernet connectivity-fault-management mep-statistics remote-mep (CIR and EIR counters enabled)**

user@host> **show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma-1 local-mep 3 remote-mep 103 count 3 remote-mep 101**

```

MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                                     :7762
CCMs received out of sequence                : 0
LBMs sent                                    : 0
Valid in-order LBRs received                 : 0
Valid out-of-order LBRs received             : 0
LBRs received with corrupted data            : 0
LBRs sent                                    : 0
LTMs sent                                    : 0
LTMs received                                : 0
LTRs sent                                    : 0
LTRs received                                : 0
Sequence number of next LTM request          : 0
1DMs sent                                    : 5
Valid 1DMs received                          : 0
Invalid 1DMs received                        : 0
DMMs sent                                    : 5
DMRs sent                                    : 0
Valid DMRs received                          : 5
Invalid DMRs received                        : 0
LMM sent                                     : 5
Valid LMM received                           : 5
Invalid LMM received                         : 0
LMR sent                                     : 0

```

```

Valid LMR received           : 5
Invalid LMR received         : 0
Remote MEP identifier         : 101
Remote MAC address           : 00:05:85:73:39:4a

```

#### Delay measurement statistics:

| Index | One-way delay<br>(usec) | Two-way delay<br>(usec) |
|-------|-------------------------|-------------------------|
| 1     | 259                     | 519                     |
| 2     | 273                     | 550                     |
| 3     | 287                     | 571                     |
| 4     | 299                     | 610                     |
| 5     | 313                     | 650                     |

```

Average one-way delay           : 286 usec
Average one-way delay variation : 62 usec
Best case one-way delay         : 259 usec
Average two-way delay           : 580 usec
Average two-way delay variation : 26 usec
Best case two-way delay         : 519 usec
Worst case two-way delay        : 650 usec

```

#### Loss measurement statistics:

| Index | Near-end<br>Frame loss<br>(CIR) | Far-end<br>Frame loss<br>(CIR) | Near-end<br>Frame loss<br>(EIR) | Far-end<br>Frame loss<br>(EIR) |
|-------|---------------------------------|--------------------------------|---------------------------------|--------------------------------|
| 1     | 10                              | 8                              | 5                               | 12                             |
| 2     | 12                              | 7                              | 6                               | 16                             |
| 3     | 7                               | 5                              | 0                               | 2                              |
| 4     | 9                               | 6                              | 8                               | 2                              |
| 5     | 3                               | 6                              | 6                               | 4                              |

```

Average near-end loss (CIR)           : 6.2
Average near-end loss ratio (CIR)      : 6.2%
Average far-end loss (CIR)            : 6.2
Average far-end loss ratio (CIR)       : 6.2%
Near-end best case loss (CIR)          : 3
Near-end best case loss ratio (CIR)    : 3%
Near-end worst case loss (CIR)         : 9
Near-end worst case loss ratio (CIR)   : 9%
Far-end best case loss (CIR)           : 5
Far-end best case loss ratio (CIR)     : 5%
Far-end worst case loss (CIR)          : 9
Far-end worst case loss ratio (CIR)    : 9%

```

|                                      |        |
|--------------------------------------|--------|
| Average near-end loss (EIR)          | : 4    |
| Average near-end loss ratio (EIR)    | : 4%   |
| Average far-end loss (EIR)           | : 3.4  |
| Average far-end loss ratio (EIR)     | : 3.4% |
| Near-end best case loss (EIR)        | : 0    |
| Near-end best case loss ratio (EIR)  | : 0%   |
| Near-end worst case loss (EIR)       | : 8    |
| Near-end worst case loss ratio (EIR) | : 8%   |
| Far-end best case loss (EIR)         | : 2    |
| Far-end best case loss ratio (EIR)   | : 2%   |
| Far-end worst case loss (EIR)        | : 6    |
| Far-end worst case loss ratio (EIR)  | : 6%   |

### show oam ethernet connectivity-fault-management mep-statistics

user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma-1

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 1
  CCMs sent                      : 6550
  CCMs received out of sequence  : 0
  LBMs sent                      : 0
  Valid in-order LBRs received   : 0
  Valid out-of-order LBRs received : 0
  LBRs received with corrupted data : 0
  LBRs sent                     : 0
  LTMs sent                     : 0
  LTMs received                 : 0
  LTRs sent                     : 0
  LTRs received                 : 0
  Sequence number of next LTM request : 0
  1DMs sent                     : 5
  Valid 1DMs received           : 0
  Invalid 1DMs received         : 0
  DMMS sent                     : 5
  DMRs sent                     : 0
  Valid DMRs received           : 5
  Invalid DMRs received         : 0
  SLM sent                      : 10
  Valid SLM received            : 20
  Invalid SLM received          : 0
  SLR sent                      : 20
```



```

Valid SLR received                : 10
Invalid SLR received              : 0

Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)         (usec)
  1      259          519
  2      273          550
  3      287          571
  4      299          610
  5      313          650
Average one-way delay              : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay            : 259 usec
Worst case one-way delay           : 313 usec
Average two-way delay              : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay            : 519 usec
Worst case two-way delay           : 650 usec
Synthetic Loss measurement
statistics:
  SLM packets sent                : 100
  SLM packets received            : 0
  SLR packets sent                : 100
  SLR packets received            : 0
Accumulated SLM statistics:
Local TXFC1 value                 : 100
Local RXFC1 value                 : 100
Last Received SLR frame TXFCftc   : 100
Last Received SLR frame TXFCbtc   : 100
SLM Frame Loss:
Frame Loss (far-end)              : 0 (0.00 %)
Frame Loss (near-end)             : 0 (0.00 %)

```

**show oam ethernet connectivity-fault-management mep-statistics remote-mep**

**user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md1 maintenance-association ma1 remote-mep 101**

```

MEP identifier: 100, MAC address: 00:05:85:73:7b:39
CCMs sent                : 7762
CCMs received out of sequence : 0

```

```

LBMs sent : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
lDMs sent : 5
Valid lDMs received : 0
Invalid lDMs received : 0
DMMs sent : 5
DMRs sent : 0
Valid DMRs received : 5
Invalid DMRs received : 0
SLM sent : 10
Valid SLM received : 20
Invalid SLM received : 0
SLR sent : 20
Valid SLR received : 10
Invalid SLR received : 0

Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      259        519
    2      273        550
    3      287        571
    4      299        610
    5      313        650
Average one-way delay : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay : 259 usec
Worst case one-way delay : 313 usec
Average two-way delay : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay : 519 usec
Worst case two-way delay : 650 usec
statistics:
SLM packets sent : 100
Synthetic Loss measurement

```

```

SLM packets received          : 0
SLR packets sent              : 100
SLR packets received          : 0
Accumulated SLM statistics:
Local TXFC1 value             : 100
Local RXFC1 value             : 100
Last Received SLR frame TXFCftc : 100
Last Received SLR frame TXFCbtc : 100
SLM Frame Loss:
Frame Loss (far-end)          : 0 (0.00 %)
Frame Loss (near-end)         : 0 (0.00 %)

```

**show oam ethernet connectivity-fault-management mep-statistics local-mep remote-mep**

**user@host> show oam ethernet connectivity-fault-management mep-statistics maintenance-domain  
md1 maintenance-association ma1 local-mep 121 remote-mep 101**

```

MEP identifier: 1, MAC address: 3c:61:04:25:fa:95
Remote MEP count: 1
  CCMs sent                      : 875
  CCMs Received                  : 300
  CCMs received out of sequence  : 0
  LBMs sent                      : 0
  Valid in-order LBRs received   : 0
  Valid out-of-order LBRs received : 0
  LBRs received with corrupted data : 0
  LBRs sent                      : 0
  LTMs sent                      : 0
  LTMs received                  : 0
  LTRs sent                      : 0
  LTRs received                  : 0
  Sequence number of next LTM request : 12
  lDMs sent                      : 0
  Valid lDMs received            : 0
  Invalid lDMs received          : 0
  Out of sync lDMs received      : 0
  DMMS sent                      : 0
  Valid DMMS received            : 0
  Invalid DMMS received          : 0
  DMRs sent                      : 0
  Valid DMRs received            : 0
  Invalid DMRs received          : 0
  LMMS sent                      : 0

```

|                       |     |
|-----------------------|-----|
| Valid LMMS received   | : 0 |
| Invalid LMMS received | : 0 |
| LMRs sent             | : 0 |
| Valid LMRs received   | : 0 |
| Invalid LMRs received | : 0 |
| SLMs sent             | : 0 |
| Valid SLMs received   | : 0 |
| Invalid SLMs received | : 0 |
| SLRs sent             | : 0 |
| Valid SLRs received   | : 0 |
| Invalid SLRs received | : 0 |
| Valid AISS generated  | : 0 |
| Valid AISS received   | : 0 |

## show oam ethernet connectivity-fault-management path-database

### Syntax

```
show oam ethernet connectivity-fault-management path-database
<host-mac-address>
<maintenance-association ma-name>
<maintenance-domain domain-name>
```

### Release Information

Command introduced in Junos OS Release 8.4.

### Description

On M7i and M10i with Enhanced CFEB (CFEB-E), M320, MX Series, ACX Series, T320, and T640 routers, display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management path database information for a host configured with an MEP.

### Options

**host-mac-address**—(Optional) Display connectivity fault management path database information for a specified Ethernet host.

**maintenance-association *ma-name***—(Optional) Display connectivity fault management path database information for the specified maintenance association.

**maintenance-domain *domain-name***—(Optional) Display connectivity fault management path database information for the specified maintenance domain.

### Required Privilege Level

view

### List of Sample Output

[show oam ethernet connectivity-fault-management path-database on page 1933](#)

### Output Fields

[Table 170 on page 1932](#) lists the output fields for the **show oam ethernet connectivity-fault-management path-database** command. Output fields are listed in the approximate order in which they appear.

Table 170: show oam ethernet connectivity-fault-management path-database Output Fields

| Field Name   | Field Description                           |
|--------------|---------------------------------------------|
| Linktrace to | MAC address of the remote MEPs in the path. |
| Interface    | Interface identifier.                       |

Table 170: show oam ethernet connectivity-fault-management path-database Output Fields (*continued*)

| Field Name                          | Field Description                          |
|-------------------------------------|--------------------------------------------|
| <b>Maintenance domain name</b>      | Maintenance domain name.                   |
| <b>Format (Maintenance domain)</b>  | Maintenance domain name format configured. |
| <b>Level</b>                        | Maintenance domain level configured.       |
| <b>Maintenance association name</b> | Maintenance association name.              |
| <b>Local Mep</b>                    | Local MEP identifier.                      |

## Sample Output

```
show oam ethernet
connectivity-fault
-management
path-database
```

```
user@host> show oam ethernet connectivity-fault-management path-database maintenance-domain
md1 maintenance-association ma1 00:05:85:79:39:ef
```

```
Linktrace to 00:05:85:79:39:ef, Interface : ge-3/0/0
Maintenance Domain: mdl, Level: 7
Maintenance Association: ma1, Local Mep: 201
```

## show oam ethernet connectivity-fault-management policer

### Syntax

```
show oam ethernet connectivity-fault-management policer <maintenance-domain md-name>
<maintenance-association ma-name>
```

### Release Information

Command introduced in Junos OS Release 10.0.

### Description

On M120, M320, MX Series, T320, and T640 routers displays connectivity-fault-management policer statistics.

### Options

This command has the following options:

**maintenance-domain *md-name***—Name of an existing CFM maintenance domain. If this option is not specified, policer statistics are displayed for all maintenance associations for all maintenance domains.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association. If this option is not specified, policer statistics are displayed for all maintenance associations for given maintenance domain. This option cannot be specified without specifying maintenance-domain name.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[clear oam ethernet connectivity-fault-management policer](#) | [1442](#)

### List of Sample Output

[show oam ethernet connectivity-fault-management policer on page 1935](#)

[show oam ethernet connectivity-fault-management policer maintenance-domain \*md-name\* on page 1936](#)

[show oam ethernet connectivity-fault-management policer maintenance-domain \*md-name\* maintenance-association \*ma-name\* on page 1936](#)

### Output Fields

[Table 171 on page 1935](#) lists the output fields for the **show oam ethernet connectivity-fault-management policer** command. Output fields are listed in the approximate order in which they appear.

Table 171: show oam ethernet connectivity-fault-management policer Output Fields

| Field Name                     | Field Description                                                                                                                                                                                                                                                              |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Legend for Policer</b>      | Describes the symbols used under the <b>Scope</b> and <b>Type</b> headings: <ul style="list-style-type: none"> <li>• G - Global scope</li> <li>• S - Service scope</li> <li>• cc - Continuity check (Type)</li> </ul>                                                          |
| <b>Maintenance Domain</b>      | Displays the maintenance domain name.                                                                                                                                                                                                                                          |
| <b>Level</b>                   | Displays the maintenance domain level configured.                                                                                                                                                                                                                              |
| <b>Maintenance association</b> | Displays the maintenance association name.                                                                                                                                                                                                                                     |
| <b>Policer</b>                 | Displays the policer name.                                                                                                                                                                                                                                                     |
| <b>Type</b>                    | Policer type. Value <b>cc</b> means this policer is used only to police continuity check CFM messages. Value <b>other</b> means this policer is used only to police non-continuity check CFM messages. Value <b>all</b> means this policer is used to police all CFM messages. |
| <b>Scope</b>                   | Policer scope. Displays whether the <i>global</i> (G) policer configuration is applicable or the session (S) specific policer config is applicable.                                                                                                                            |
| <b>Drop count</b>              | Displays the number of packets dropped by the indicated policer.                                                                                                                                                                                                               |

## Sample Output

```
show oam ethernet
connectivity-fault
-management
policer
```

Displays the policer information for all maintenance associations and their maintenance domains.

```
show oam ethernet connectivity-fault-management policer
```

```
Legend for Policer
G - Global scope
S - Service scope
cc - Continuity check
```



```

Maintenance Domain: md1 Level: 1
Maintenance association Policer      Type      Scope Drop count
ma1                      cfm-policer1 all       G          300
ma1-2                    cfm-policer1 cc        S          259
ma1-2                    cfm-policer1 other    G          300
Maintenance Domain: md2 Level: 2
Maintenance association Policer      Type      Scope Drop count
ma2                      cfm-policer1 cc        G          300
ma2                      cfm-policer2 other    S          223

```

**show oam ethernet  
connectivity-fault  
-management  
policer  
maintenance-domain  
md-name**

Displays the policer information for the specified maintenance domain and its maintenance associations.

```
show oam ethernet connectivity-fault-management policer maintenance-domain
md1
```

```

Legend for Policer
G - Global scope
S - Service scope
cc - Continuity check

Maintenance Domain: md1 Level: 1
Maintenance association Policer      Type      Scope Drop count
ma1                      cfm-policer1 all       G          300
ma1-2                    cfm-policer1 cc        S          259
ma1-2                    cfm-policer1 other    G          300

```

**show oam ethernet  
connectivity-fault  
-management  
policer  
maintenance-domain  
md-name  
maintenance-association  
ma-name**

Displays the policer information for the specified **maintenance-domain** *md-name* and **maintenance-association** *ma-name*.

```
show oam ethernet connectivity-fault-management policer maintenance-domain  
md5 maintenance-association ma5
```

Legend for Policer

G - Global scope

S - Service scope

cc - Continuity check

Maintenance Domain: md5 Level: 5

| Maintenance association | Policer       | Type  | Scope | Drop count |
|-------------------------|---------------|-------|-------|------------|
| ma5                     | cfm-policer   | cc    | S     | 187        |
| ma5                     | cfm-policer-2 | other | S     | 234        |

## show oam ethernet connectivity-fault-management sla-iterator-statistics

### Syntax

```
show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md-name
maintenance-association ma-name
sla-iterator sla-iterator
<local-mep local-mep-id>
<remote-mep remote-mep-id>
```

### Release Information

Command introduced in Junos OS Release 11.4 for EX Series switches.

Command introduced in Junos OS Release 9.6.

Command introduced in Junos OS Release 12.2 for ACX Series routers.

Command introduced in Junos OS Release 13.2 for MX Series routers (not on MPC3E Hyperion cards).

### Description

Display the Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics.

### Options

**maintenance-domain *md-name***—Name of an existing connectivity fault management (CFM) maintenance domain.

**maintenance-association *ma-name***—Name of an existing CFM maintenance association.

**sla-iterator *sla-iterator***— Name of the iterator profile.

**local-mep *local-mep-id***—(Optional) Numeric identifier of the local MEP. The range of values is **1** through **8191**.

**remote-mep *remote-mep-id***—(Optional) Numeric identifier of the remote MEP. The range of values is **1** through **8192**.

### Required Privilege Level

view

### RELATED DOCUMENTATION

*Configuring an Iterator Profile on a Switch (CLI Procedure)*

*clear oam ethernet connectivity-fault-management sla-iterator-statistics*

### List of Sample Output

[show oam ethernet connectivity-fault-management sla-iterator-statistics](#) on page 1942

[show oam ethernet connectivity-fault-management sla-iterator-statistics \(MX Series routers\)](#) on page 1942

[show oam ethernet connectivity-fault-management sla-iterator-statistics \(MX Series routers\) - Delay Measurement \(DM\)in Metro Ethernet Forum \(MEF\) mode](#) on page 1944

[show oam ethernet connectivity-fault-management sla-iterator-statistics \(MX Series routers\) - Synthetic loss measurement \(SLM\) in Metro Ethernet Forum \(MEF\) mode](#) on page 1945

[show oam ethernet connectivity-fault-management sla-iterator-statistics \(MX Series routers\) - Delay measurement \(DM\) statistics in non-Metro Ethernet Forum \(MEF\) mode](#) on page 1947

[show oam ethernet connectivity-fault-management sla-iterator-statistics \(MX Series routers\) - Synthetic loss measurement \(SLM\) statistics in non-Metro Ethernet Forum \(MEF\) mode](#) on page 1947

[show oam ethernet connectivity-fault-management sla-iterator-statistics \(MX Series routers\) - Delay Measurement \(DM\) with “legacy-pm-display” option in enhanced-cfm mode](#) on page 1948

[show oam ethernet connectivity-fault-management state](#) - To verify the Connectivity Fault Management (CFM) state on page 1949

## Output Fields

[Table 172 on page 1939](#) lists the output fields for the **show oam ethernet connectivity-fault-management sla-iterator-statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 172: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields**

| Output Field Name              | Output Field Description                                                                                     |
|--------------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>Maintenance domain</b>      | Name of the maintenance domain.                                                                              |
| <b>Level</b>                   | Level of the maintenance domain level configured.                                                            |
| <b>Maintenance association</b> | Name of the maintenance association.                                                                         |
| <b>Local MEP id</b>            | Numeric identifier of the local MEP.                                                                         |
| <b>Remote MEP id</b>           | Numeric identifier of the remote MEP.                                                                        |
| <b>Remote MAC address</b>      | Unicast MAC address of the remote MEP.                                                                       |
| <b>Iterator name</b>           | Name of iterator.                                                                                            |
| <b>Iterator Id</b>             | Numeric identifier of the iterator.                                                                          |
| <b>Iterator cycle time</b>     | Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection |
| <b>Iteration period</b>        | Maximum number of cycles per iteration                                                                       |
| <b>Iterator status</b>         | Current status of iterator whether running or stopped.                                                       |

Table 172: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)

| Output Field Name                               | Output Field Description                                                                                |
|-------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| <b>Infinite iterations</b>                      | Status of iteration as infinite or finite.                                                              |
| <b>Counter reset time</b>                       | Date and time when the counter was reset.                                                               |
| <b>Reset reason</b>                             | Reason to reset counter.                                                                                |
| <b>Delay weight</b>                             | Calculation weight of delay.                                                                            |
| <b>Delay variation weight</b>                   | Calculation weight of delay variation.                                                                  |
| <b>DMM sent</b>                                 | Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.                        |
| <b>DMM skipped for threshold hit</b>            | Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.                 |
| <b>DMM skipped for threshold hit window</b>     | Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window. |
| <b>DMR received</b>                             | Number of delay measurement reply (DMR) frames received.                                                |
| <b>DMR out of sequence</b>                      | Total number of DMR out of sequence packets received.                                                   |
| <b>DMR received with invalid time stamps</b>    | Total number of DMR frames received with invalid timestamps.                                            |
| <b>Average two-way delay</b>                    | Average two-way frame delay for the statistics displayed.                                               |
| <b>Average two-way delay variation</b>          | Average two-way "frame jitter" for the statistics displayed.                                            |
| <b>Average one-way forward delay variation</b>  | Average one-way forward delay variation for the statistics displayed in microseconds.                   |
| <b>Average one-way backward delay variation</b> | Average one-way backward delay variation for the statistics displayed in microseconds.                  |
| <b>Weighted average two-way delay</b>           | Weighted average two-way delay for the statistics displayed in microseconds.                            |
| <b>Weighted average two-way delay variation</b> | Weighted average two-way delay variation for the statistics displayed in microseconds.                  |

Table 172: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)

| Output Field Name                                        | Output Field Description                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Weighted average one-way backward delay variation</b> | Weighted average one-way backward delay variation for the statistics displayed in microseconds.                                                                                                                                                                                                                                   |
| <b>Weighted average one-way forward delay variation</b>  | Weighted average one-way forward delay variation for the statistics displayed in microseconds.                                                                                                                                                                                                                                    |
| <b>SLM packets sent</b>                                  | Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.                                                                                                                                                                                                   |
| <b>SLM packets received</b>                              | Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.                                                                                                                                                                                             |
| <b>SLR packets sent</b>                                  | Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.                                                                                                                                                                                                 |
| <b>SLR packets received</b>                              | Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.                                                                                                                                                                                           |
| <b>Local TXFC1 value</b>                                 | Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.   |
| <b>Local RXFC1 value</b>                                 | Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet. |
| <b>Last Received SLR frame TXFCf(tc)</b>                 | Value of the local counter TxFCI at the time of SLM frame transmission.                                                                                                                                                                                                                                                           |
| <b>Last Received SLR frame TXFCb(t)</b>                  | Value of the local counter RxFCI at the time of SLR frame transmission.                                                                                                                                                                                                                                                           |
| <b>Frame loss (near-end)</b>                             | Count of frame loss associated with ingress data frames.                                                                                                                                                                                                                                                                          |
| <b>Frame loss (far-end)</b>                              | Count of frame loss associated with egress data frames.                                                                                                                                                                                                                                                                           |

## Sample Output

**show oam ethernet connectivity-fault  
-management sla-iterator-statistics**

**user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics sla-iterator  
i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1 remote-mep 2**

```

Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i1, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent : 23898520
DMM skipped for threshold hit : 11000
DMM skipped for threshold hit window : 0
DMR received : 23851165
DMR out of sequence : 1142
DMR received with invalid time stamps : 36540
Average two-way delay : 129 usec
Average two-way delay variation : 15 usec
Average one-way forward delay variation : 22 usec
Average one-way backward delay variation : 22 usec
Weighted average two-way delay : 134 usec
Weighted average two-way delay variation : 8 usec
Weighted average one-way forward delay variation : 6 usec
Weighted average one-way backward delay variation : 2 usec

```

## Sample Output

**show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers)**

**user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics  
maintenance-domain mdu maintenance-association mau local-mep 4 remote-mep 3 sla-iterator lm**

## Iterator statistics:

Maintenance domain: 2, Level: 2

Maintenance association: W-160432000-001, Local MEP id: 2

Remote MEP id: 1, Remote MAC address: 00:90:69:0a:43:39

Iterator name: iter1, Iterator Id: 1

Iterator cycle time: 100ms, Iteration period: 10 cycles

Iterator status: running, Infinite iterations: true

Counter reset time: 2012-09-25 02:15:31 PDT (00:00:45 ago)

Reset reason: Adjacency flap

## Iterator loss measurement statistics:

LMM sent : 444

LMM skipped for threshold hit : 0

LMM skipped for threshold hit window: 0

LMR received : 444

LMR out of sequence : 0

LMR forwarding-class mismatch : 0

## Accumulated transmit statistics:

Near-end (CIR) : 0

Far-end (CIR) : 0

Near-end (EIR) : 0

Far-end (EIR) : 0

## Accumulated receive statistics:

Near-end (CIR) : 0

Far-end (CIR) : 0

Near-end (EIR) : 0

Far-end (EIR) : 0

## Accumulated loss statistics:

Near-end loss (CIR) : 0

Near-end loss-ratio (CIR) : 0 (0.00000%)

Far-end loss (CIR) : 0

Far-end loss-ratio (CIR) : 0 (0.00000%)

Near-end loss (EIR) : 0

Near-end loss-ratio (EIR) : 0 (0.00000%)

Far-end loss (EIR) : 0

Far-end loss-ratio (EIR) : 0 (0.00000%)

## Last loss measurement statistics:

Near-end (CIR) : 0

Far-end (CIR) : 0

Near-end (EIR) : 0

Far-end (EIR) : 0



## Sample Output

**show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Delay Measurement (DM)in Metro Ethernet Forum (MEF) mode**

**user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics**  
**maintenance-domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator DM**

```

Iterator statistics:
  Maintenance domain: md6-1, Level: 6
  Maintenance association: ma1, Local MEP id: 1
  Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
  Iterator name: Test_DM, Iterator Id: 1
  Iterator cycle time: 1000ms, Iteration period: 2000 cycles
  Iterator status: running, Infinite iterations: true
  Counter reset time: 2018-12-05 19:48:23 PST (00:00:54 ago)
  Reset reason: Adjacency flap
  Current delay measurement statistics:
    Measurement Interval Index          : 2 (Suspect status : 1)
    Measurement Interval Start Time      : 2018-12-05 19:48:23 PST (Elapsed
time : 53011 msec)
    Frame Delay two way (min, max, avg)   : 251 , 295 , 262 (usec)
    Frame Delay forward (min, max, avg)   : 125 , 147 , 131 (usec)
    Frame Delay backward (min, max, avg)   : 125 , 147 , 131 (usec)
    Inter Frame Delay two way (min, max, avg) : 0 , 42 , 7 (usec)
    Inter Frame Delay forward (min, max, avg) : 0 , 21 , 3 (usec)
    Inter Frame Delay backward (min, max, avg) : 0 , 21 , 3 (usec)
    Frame Delay Range two way (max, avg)    : 0 , 0 (usec)
    Frame Delay Range forward (max, avg)    : 0 , 0 (usec)
    Frame Delay Range backward (max, avg)   : 0 , 0 (usec)
    SOAM TXed                             : 52
    SOAM RXed                             : 52

Delay measurement bin statistics:
  Measurement Interval Index          : 2
  Two Way Frame Delay
    [0 - 4999] (usec) : 52
    [5000 - 9999] (usec) : 0
    [10000 - Infinity] (usec) : 0
  Forward Frame Delay
    [0 - 4999] (usec) : 52
    [5000 - 9999] (usec) : 0
    [10000 - Infinity] (usec) : 0
  Backward Frame Delay
    [0 - 4999] (usec) : 52

```

```

[5000      - 9999      ] (usec)   : 0
[10000     - Infinity  ] (usec)   : 0
Two Way Inter Frame Delay Variation
[0          - 4999      ] (usec)   : 51
[5000      - 9999      ] (usec)   : 0
[10000     - Infinity  ] (usec)   : 0
Forward Inter Frame Delay Variation
[0          - 4999      ] (usec)   : 51
[5000      - 9999      ] (usec)   : 0
[10000     - Infinity  ] (usec)   : 0
Backward Inter Frame Delay Variation
[0          - 4999      ] (usec)   : 51
[5000      - 9999      ] (usec)   : 0
[10000     - Infinity  ] (usec)   : 0
Two Way Frame Delay Range
[0          - 4999      ] (usec)   : 0
[5000      - 9999      ] (usec)   : 0
[10000     - Infinity  ] (usec)   : 0
Forward Frame Delay Range
[0          - 4999      ] (usec)   : 0
[5000      - 9999      ] (usec)   : 0
[10000     - Infinity  ] (usec)   : 0
Backward Frame Delay Range
[0          - 4999      ] (usec)   : 0
[5000      - 9999      ] (usec)   : 0
[10000     - Infinity  ] (usec)   : 0

```

**show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Synthetic loss measurement (SLM) in Metro Ethernet Forum (MEF) mode**

user@router> **show oam ethernet connectivity-fault-management sla-iterator-statistics**

**maintenance-domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator SLM**

```

Iterator statistics:
Maintenance domain: md6-1, Level: 6
Maintenance association: ma1, Local MEP id: 1
Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
Iterator name: Test_SLM, Iterator Id: 2
Iterator cycle time: 1000ms, Iteration period: 2000 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2018-12-05 19:48:23 PST (00:00:11 ago)
Reset reason: Adjacency flap
Iterator synthetic loss measurement statistics:

```

```

SLM skipped for threshold hit      : 0
SLM skipped for threshold hit window : 0
SLR out of sequence                : 0
SLM Sample Size                    : 180 SLMs
Local packet Stats (TxFCI, RxFCI)  : 11, 11
Last SLR packet stats (TxFCf, TxFCb) : 11, 11
Last measured FLR (fwd, bkwd)      : 0.000%, 0.000% (Sample #NA)

```

Current Measurement Interval loss statistics:

```

Measurement Interval Index          : 2 (Suspect Status: 1)
Measurement Interval Start Time      : 2018-12-05 19:48:23 PST (Elapsed time:
10042 msec)
SOAM Frames (Tx, Rx)                : 10, 10
Forward Frame Stats (Tx, Rx)         : 10, 10
Backward Frame Stats (Tx, Rx)        : 10, 10
Frame Loss (fwd, bkwd)              : 0, 0
Forward FLR minimum                  : 0.000%
Forward FLR maximum                  : 0.000%
Forward FLR average                  : 0.000%
Backward FLR minimum                 : 0.000%
Backward FLR maximum                 : 0.000%
Backward FLR average                 : 0.000%

```

Current Measurement Interval availability statistics:

```

Measurement Interval Index          : 2 (Suspect Status: 1)
Measurement Interval Start Time      : 2018-12-05 19:48:23 PST (Elapsed time:
10042 msec)
High loss (fwd, bkwd)               : 0 , 0
Consecutive high loss (fwd, bkwd)   : 0 , 0
Available (fwd, bkwd)               : 0 , 0
Unavailable (fwd, bkwd)             : 0 , 0
Forward FLR minimum                  : 0.000%
Forward FLR maximum                  : 0.000%
Forward FLR average                  : 0.000%
Backward FLR minimum                 : 0.000%
Backward FLR maximum                 : 0.000%
Backward FLR average                 : 0.000%

```

```

Last known available status (fwd, bkwd) : unknown, unknown
Last known forward availability transition : NA
Last known backward availability transition : NA

```

# show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Delay measurement (DM) statistics in non-Metro Ethernet Forum (MEF) mode

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator DM
```

```
Iterator statistics:
Maintenance domain: md6-1, Level: 6
Maintenance association: ma1, Local MEP id: 1
Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
Iterator name: Test_DM, Iterator Id: 1
Iterator cycle time: 1000ms, Iteration period: 2000 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2018-12-05 19:43:47 PST (00:02:44 ago)
Reset reason: Adjacency flap
Iterator delay measurement statistics:
Calculation weight: Delay: 1, Delay variation: 1
DMM sent : 164
DMM skipped for threshold hit : 0
DMM skipped for threshold hit window : 0
DMR received : 164
DMR out of sequence : 0
DMR forwarding-class mismatch : 0
DMR received with invalid time stamps : 0
Average two-way delay : 234 usec
Average two-way delay variation : 9 usec
Average one-way forward delay variation : 346 usec
Average one-way backward delay variation : 346 usec
Weighted average two-way delay : 221 usec
Weighted average two-way delay variation : 2 usec
Weighted average one-way forward delay variation : 357 usec
Weighted average one-way backward delay variation: 355 usec
Bestcase two-way delay : 210 usec
Worstcase two-way delay : 283 usec
Weighted Bestcase two-way delay : 210 usec
Weighted Worstcase two-way delay : 283 usec
```

# show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Synthetic loss measurement (SLM) statistics in non-Metro Ethernet Forum (MEF) mode

```
user@router> show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator SLM
```

```
Iterator statistics:
```

```

Maintenance domain: md6-1, Level: 6
Maintenance association: ma1, Local MEP id: 1
Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
Iterator name: Test_SLM, Iterator Id: 2
Iterator cycle time: 1000ms, Iteration period: 2000 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2018-12-05 19:43:47 PST (00:03:23 ago)
Reset reason: Adjacency flap
Iterator synthetic loss measurement statistics:
    SLM sent : 202
    SLM skipped for threshold hit : 0
    SLM skipped for threshold hit window: 0
    SLR received : 202
    SLR out of sequence : 0
SLM transmit statistics:
    SLM TXed : 202
    SLM RXed : 202
    SLM Last packet Tx count : 202
    SLM Last packet Rx count : 202
Last loss measurement statistics:
    Near-end loss : 0 (0.00%)
    Far-end loss : 0 (0.00%)

```

**show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) - Delay Measurement (DM) with “legacy-pm-display” option in enhanced-cfm mode**

user@router> **show oam ethernet connectivity-fault-management sla-iterator-statistics**  
**maintenance-domain md6-1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator Legacy**

```

Iterator statistics:
Maintenance domain: md6-1, Level: 6
Maintenance association: ma1, Local MEP id: 1
Remote MEP id: 2, Remote MAC address: 00:23:9c:db:0d:7a
Iterator name: Test_DM_Legacy, Iterator Id: 3
Iterator cycle time: 1000ms, Iteration period: 2000 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2018-12-05 19:51:33 PST (00:00:15 ago)
Reset reason: Adjacency flap
Iterator delay measurement statistics:
    Calculation weight: Delay: 1, Delay variation: 1
    DMM sent : 14
    DMM skipped for threshold hit : 0
    DMM skipped for threshold hit window : 0

```

```

DMR received : 14
DMR out of sequence : 0
DMR forwarding-class mismatch : 0
DMR received with invalid time stamps : 0
Average two-way delay : 324 usec
Average two-way delay variation : 7 usec
Average one-way forward delay variation : 240 usec
Average one-way backward delay variation : 240 usec
Weighted average two-way delay : 324 usec
Weighted average two-way delay variation : 6 usec
Weighted average one-way forward delay variation : 222 usec
Weighted average one-way backward delay variation: 222 usec
Bestcase two-way delay : 312 usec
Worstcase two-way delay : 333 usec
Weighted Bestcase two-way delay : 312 usec
Weighted Worstcase two-way delay : 333 usec

```

#### **show oam ethernet connectivity-fault-management state - To verify the Connectivity Fault Management (CFM) state**

The following command is to verify whether the CFM state is in enhanced-cfm mode or not.

**user@router> show oam ethernet connectivity-fault-management state**

```

Connectivity fault management state:
  CFM Mode Of Operation: Enhanced
  Enhanced IP Mode: Enabled
  CFM Config State: Ok
  CFM Cleanup State: Ok
  CFM Restart Timer State: Cleanup Timer State stopped Rebooting in 0 sec
  CFM CFMMAN Job State: Not Pending
  Number of sessions: 1
  Number of sessions created: 1
  Number of sessions deleted: 0
  Number of sessions freed: 0
  Number of sessions enqueued: 1
  Number of sessions dequeued: 1
  VPLS feature: enabled
  Token based forwarding feature: enabled
  Forwarding table filtering simulation feature: disabled
  Hardware assisted flooding feature: enabled
  Flood resynchronization for GRES feature: enabled
  Shared interface filter feature: disabled
  Hardware timestamping feature: disabled

```

```
Marking of connection protection TLV feature: disabled
CFMD config memory resource  limit(in bytes): 3221225472
CFMD max resident set (peak) size (in bytes):  24158208
```

Packet processing state:

```
State of the connection to packet processing daemon: down
State of the flow to packet processing daemon: ready
State of the packet processing job: ready
Number of times the connection to packet processing daemon was blocked: 0
State of the connection to cfmmman: slots: 2 3 4 5
```

Filter state:

```
State of the connection to firewall daemon: Connected
Number of reconnects made to firewall daemon: 0
Number of requests sent to firewall daemon: 13
Number of requests accepted by firewall daemon: 13
Number of requests rejected by firewall daemon: 0
Number of requests lost due to disconnection: 0
```

## show oam ethernet connectivity-fault-management synthetic-loss-statistics

### Syntax

```
show oam ethernet connectivity-fault-management synthetic-loss-statistics
<local-mep local-mep-id>
maintenance-association ma-name
<count entry-count>
maintenance-domain md-name
<remote-mep remote-mep-id>
```

### Release Information

Command introduced in Junos OS Release 13.2 for MX Series routers.

### Description

On MX Series routers with Modular Port Concentrators (MPCs) with Ethernet interfaces, display the on-demand ETH-SLM statistics.

### Options

**count entry-count**—(Optional) Number of entries to display from the statistics table. The range of values is from 1 through 100. The default value is 100.

**local-mep local-mep-id**—(Optional) Numeric identifier of the local MEP. The range of values is from 1 through 8192.

**maintenance-association ma-name**—Name of an existing CFM maintenance association.

**maintenance-domain md-name**—Name of an existing connectivity fault management (CFM) maintenance domain.

**remote-mep remote-mep-id**—(Optional) Numeric identifier of the remote MEP. The range of values is from 1 through 8192.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[clear oam ethernet connectivity-fault-management statistics](#) | 1444

[clear oam ethernet connectivity-fault-management synthetic-loss-measurement](#) | 1446

[show oam ethernet connectivity-fault-management interfaces](#) | 1879

[show oam ethernet connectivity-fault-management mep-database](#) | 1901

[show oam ethernet connectivity-fault-management mep-statistics](#) | 1916



## List of Sample Output

[show oam ethernet connectivity-fault-management synthetic-loss-statistics on page 1953](#)

### Output Fields

[Table 173 on page 1952](#) lists the output fields for the **show oam ethernet connectivity-fault-management synthetic-loss-statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 173: show oam ethernet connectivity-fault-management synthetic-loss-statistics Output Fields**

| Output Field Name                        | Field Description                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MEP identifier</b>                    | Maintenance association end point (MEP) numeric identifier.                                                                                                                                                                                                                                                                       |
| <b>MAC address</b>                       | Unicast MAC address configured for the MEP.                                                                                                                                                                                                                                                                                       |
| <b>Remote MEP count</b>                  | Number of remote MEPs (unless you specify the <b>remote-mep</b> option).                                                                                                                                                                                                                                                          |
| <b>Remote MEP identifier</b>             | Numeric identifier of the remote MEP.                                                                                                                                                                                                                                                                                             |
| <b>Remote MAC address</b>                | Unicast MAC address of the remote MEP.                                                                                                                                                                                                                                                                                            |
| <b>SLM packets sent</b>                  | Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.                                                                                                                                                                                                   |
| <b>SLM packets received</b>              | Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.                                                                                                                                                                                             |
| <b>SLR packets sent</b>                  | Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.                                                                                                                                                                                                 |
| <b>SLR packets received</b>              | Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.                                                                                                                                                                                           |
| <b>Local TXFC1 value</b>                 | Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.   |
| <b>Local RXFC1 value</b>                 | Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TxFC1 is sent in the packet. |
| <b>Last Received SLR frame TXFCf(tc)</b> | Value of the local counter TxFCI at the time of SLM frame transmission.                                                                                                                                                                                                                                                           |

**Table 173: show oam ethernet connectivity-fault-management synthetic-loss-statistics Output Fields (continued)**

| Output Field Name                      | Field Description                                                       |
|----------------------------------------|-------------------------------------------------------------------------|
| <b>Last Received SLR frame TXFCb(t</b> | Value of the local counter RxFCI at the time of SLR frame transmission. |
| <b>Frame loss (near-end)</b>           | Count of frame loss associated with ingress data frames.                |
| <b>Frame loss (far-end)</b>            | Count of frame loss associated with egress data frames.                 |

## Sample Output

**show oam ethernet connectivity-fault-management  
synthetic-loss-statistics**

user@switch> **show oam ethernet connectivity-fault-management synthetic-loss-statistics  
maintenance-domain md6 maintenance-association ma6**

```
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
  Remote MEP identifier: 101
  Remote MAC address: 00:05:85:73:39:4a
  Synthetic Loss measurement statistics:
    SLM packets sent           : 100
    SLM packets received       : 0
    SLR packets sent           : 100
    SLR packets received       : 0
  Accumulated SLM statistics:
    Local TXFC1 value          : 100
    Local RXFC1 value          : 100
    Last Received SLR frame TXFCftc : 100
    Last Received SLR frame TXFCbtc : 100
  SLM Frame Loss:
    Frame Loss (far-end)       : 0 (0.00 %)
    Frame Loss (near-end)      : 0 (0.00 %)
```

# show oam ethernet evc

## Syntax

```
show oam ethernet evc <evc-id>
```

## Release Information

Command introduced in Junos OS Release 9.5.

## Description

On MX Series routers with OAM Ethernet Virtual Connection (EVC) configurations, displays the EVC configuration and status information.

## Options

This command has no options.

## Required Privilege Level

View

## Output Fields

[Table 174 on page 1954](#) lists the output fields for the **show oam ethernet evc** command. Output fields are listed in the approximate order in which they appear.

Table 174: show oam ethernet evc Output Fields

| Field Name     | Field Description                                                                               |
|----------------|-------------------------------------------------------------------------------------------------|
| EVC identifier | Header for the EVC information showing the EVC name, configuration, and active/inactive status. |
| UNI count      | Number of configured and active UNIs.                                                           |
| Protocol       | Protocol configured between the UNIs.                                                           |
| Local UNIs     | Heading for the list of local UNIs                                                              |
| UNI Identifier | Name of the UNI.                                                                                |
| Interface      | Interface type-dpc/pic/port.unit-number.                                                        |
| Status         | Status operational or not operational.                                                          |

## Sample Output

**show oam ethernet evc**

user@host> **show oam ethernet evc**

```
EVC identifier: evc1, Point-to-Point, Active
UNI count: Configured(2), Active(2)
Protocol: cfm, Management domain: md, Management association: ma
Local UNIs:
  UNI Identifier      Interface      Status
  uni1                ge-1/1/1      Operational
  uni2                ge-1/1/1      Not Operational
```

# show oam ethernet fnp interface

## Syntax

```
show oam ethernet fnp interface
<ethernet-interface-name>
<routing-instance routing-instance-name>
```

## Release Information

Command introduced in Junos OS Release 11.4.

## Description

On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) information for Ethernet interfaces.

## Options

**interface-name**—(Optional) Display Ethernet FNP information for the specified Ethernet interface only.

**routing-instance-name**—(Optional) Display FNP for the specified routing instance.

## Required Privilege Level

view

## List of Sample Output

[show oam ethernet fnp interface on page 1957](#)

## Output Fields

[Table 175 on page 1956](#) lists the output fields for the **show oam ethernet fnp interface** command. Output fields are listed in the approximate order in which they appear.

Table 175: show oam ethernet fnp interface Output Fields

| Field Name             | Field Description                                    |
|------------------------|------------------------------------------------------|
| Interface              | Name of the interface for the displayed information. |
| VLAN                   | Name of the VLAN.                                    |
| State                  | Displays state of the interface.                     |
| FNP Message Interface  | Displays the message interface type.                 |
| FNP Message Source MAC | Displays the source MAC address.                     |

## Sample Output

**show oam ethernet fnp interface**

user@host> **show oam ethernet fnp interface**

| The FNP controlled interfaces are: |      |       |             |                   |  |
|------------------------------------|------|-------|-------------|-------------------|--|
| Interface                          | VLAN | State | FNP message | FNP message       |  |
|                                    |      |       | Interface   | Source MAC        |  |
| ge-0/0/0.30                        | 30   | down  | lsi.1054976 | a0:aa:aa:aa:aa:aa |  |
| ge-0/0/0.20                        | 20   | down  | lsi.1054976 | a0:aa:aa:aa:aa:aa |  |

# show oam ethernet fnp messages

## Syntax

```
show oam ethernet fnp messages
<interface interface-name>
<routing instance routing-instance-name>
```

## Release Information

Command introduced in Junos OS Release 11.4

## Description

On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) messages.

## Options

**interface-name**—(Optional) Display Ethernet FNP messages for the specified Ethernet interface only.

**routing-instance-name**—(Optional) Display FNP messages for the specified routing instance.

## Required Privilege Level

view

## List of Sample Output

[show oam ethernet fnp messages on page 1959](#)

## Output Fields

[Table 176 on page 1958](#) lists the output fields for the **show oam ethernet fnp messages** command. Output fields are listed in the approximate order in which they appear.

Table 176: show oam ethernet fnp messages Output Fields

| Field Name                      | Field Description                                                                |
|---------------------------------|----------------------------------------------------------------------------------|
| Message from source MAC address | The source MAC address of the message.                                           |
| Originating port number         | Port number of the original message.                                             |
| Time since last message         | Elapsed time in hours, minutes, and seconds since the last message was received. |
| Time since last message update  | Elapsed time in hours, minutes, and seconds since the last message was updated.  |
| Total messages received         | Number of messages received.                                                     |

Table 176: show oam ethernet fnp messages Output Fields (*continued*)

| Field Name          | Field Description                         |
|---------------------|-------------------------------------------|
| Domain ID           | Domain ID of the message.                 |
| STP Root ID         | The spanning tree Root ID of the message. |
| Trigger Reason      | The reason why the message was triggered. |
| Effectuated VLANs   | Number of VLANs that are affected.        |
| Disabled interfaces | Name of the interfaces that are disabled. |

## Sample Output

**show oam ethernet fnp messages**

user@host> **show oam ethernet fnp messages**

```

Active FNP messages on interface lsi.1054465
Message source MAC: a0:aa:aa:aa:aa:aa
  Originating port number: 141077
  Time since last message: 00:00:00
  Time since last message update: 00:00:00
  Total messages received: 1
  Domain ID: 0
  STP Root ID: 0.f0:ff:ff:ff:ff:ff
  Trigger reason: todo
  Effectuated VLANs: 10
  Disabled interfaces:
    Interface VLAN
    ge-0/0/0.10 10

```



# show oam ethernet fnp status

## Syntax

```
show oam ethernet fnp status
<interface interface-name>
<routing instance routing-instance-name>
```

## Release Information

Command introduced in Junos OS Release 11.4

## Description

On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Failure Notification Protocol (FNP) status.

## Options

**interface-name**—(Optional) Display Ethernet FNP information for the specified Ethernet interface only.

**routing-instance-name**—(Optional) Display FNP for the specified routing instance.

## Required Privilege Level

view

## List of Sample Output

[show oam ethernet fnp status on page 1961](#)

## Output Fields

[Table 177 on page 1960](#) lists the output fields for the **show oam ethernet fnp status** command. Output fields are listed in the approximate order in which they appear.

Table 177: show oam ethernet fnp status Output Fields

| Field Name             | Field Description                                                     |
|------------------------|-----------------------------------------------------------------------|
| FNP interval           | The time interval between messages.                                   |
| Loss threshold         | The number of messages that can be lost before FNP is marked as down. |
| FNP enabled interfaces | Displays interfaces that are enabled.                                 |
| Interface              | The name of the interface.                                            |
| Domain ID              | Domain ID of the message.                                             |
| STP Root ID            | The spanning tree Root ID of the message.                             |

Table 177: show oam ethernet fnp status Output Fields (*continued*)

| Field Name   | Field Description                      |
|--------------|----------------------------------------|
| FNP Messages | The total number of messages received. |

## Sample Output

**show oam ethernet fnp status**

user@host> **show oam ethernet status**

```
FNP interval:
Loss threshold
FNP enabled interfaces
Interface      Domain ID    STP Root ID      FNP Messages
ge-0/0/0.1278    100         0.f0:ff:ff:ff:ff:ff      0
```

# show oam ethernet link-fault-management

## Syntax

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

## Release Information

Command introduced in Junos OS Release 8.2.

## Description

On EX Series switches and M320, M120, MX Series, T320, and T640 routers, display Operation, Administration, and Management (OAM) link fault management information for Ethernet interfaces.

## Options

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

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

## Required Privilege Level

view

## List of Sample Output

- [show oam ethernet link-fault-management brief on page 1968](#)
- [show oam ethernet link-fault-management brief \(Loopback tracking\) on page 1968](#)
- [show oam ethernet link-fault-management detail on page 1969](#)
- [show oam ethernet link-fault-management detail \(backup Routing Engine\) on page 1970](#)

## Output Fields

[Table 178 on page 1962](#) lists the output fields for the **show oam ethernet link-fault-management** command. Output fields are listed in the approximate order in which they appear.

Table 178: show oam ethernet link-fault-management Output Fields

| Field Name | Field Description                                                                                                                                                                                                             | Level of Output |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| Status     | <p>Indicates the status of the OAM discovery state mechanism .</p> <ul style="list-style-type: none"> <li><b>Down</b>—Discovery mechanism is not running..</li> <li><b>Running</b>—Discovery mechanism is running.</li> </ul> | All levels      |

Table 178: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                    | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Level of Output |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Discovery state</b>        | <p>State of the discovery mechanism. If the status of the discovery mechanism is Down then the state of discovery mechanism is Fault. However, if the status of the discovery mechanism is Running then the state can be any one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Passive Wait</b></li> <li>• <b>Active Send Local</b></li> <li>• <b>Send Any</b></li> <li>• <b>Send Local Remote</b></li> <li>• <b>Send Local Remote Ok</b></li> <li>• <b>Fault</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | All levels      |
| <b>ISSU</b>                   | Specifies that the local end is undergoing a unified in-service software upgrade (ISSU).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Peer address</b>           | Address of the OAM peer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | All levels      |
| <b>Flags</b>                  | <p>Information about the interface. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i>.</p> <ul style="list-style-type: none"> <li>• <b>Remote-Stable</b>—Indicates remote OAM client acknowledgment of and satisfaction with local OAM state information. <b>False</b> indicates that remote DTE either has not seen or is unsatisfied with local state information. <b>True</b> indicates that remote DTE has seen and is satisfied with local state information.</li> <li>• <b>Local-Stable</b>—Indicates local OAM client acknowledgment of and satisfaction with remote OAM state information. <b>False</b> indicates that local DTE either has not seen or is unsatisfied with remote state information. <b>True</b> indicates that local DTE has seen and is satisfied with remote state information.</li> <li>• <b>Remote-State-Valid</b>—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. <b>False</b> indicates that OAM client has not seen remote state information. <b>True</b> indicates that the OAM client has seen remote state information.</li> </ul> | All levels      |
| <b>Remote loopback status</b> | Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |

Table 178: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Level of Output |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Remote entity information</b> | <p>Remote entity information.</p> <ul style="list-style-type: none"> <li>• <b>Remote MUX action</b>—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs.</li> <li>• <b>Remote parser action</b>—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs.</li> <li>• <b>Discovery mode</b>—Indicates whether discovery mode is active or inactive.</li> <li>• <b>Unidirectional mode</b>—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes.</li> <li>• <b>Remote loopback mode</b>—Indicates whether remote loopback is supported or unsupported.</li> <li>• <b>Link events</b>—Indicates whether interpreting link events is supported or unsupported on the remote peer.</li> <li>• <b>Variable requests</b>—Indicates whether variable requests are supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.</li> <li>• <b>Remote in ISSU</b>—Indicates that the remote end is undergoing a unified in-service software upgrade (ISSU).</li> </ul> | All levels      |
| <b>Loopback Tracking</b>         | Indicates that loopback detection is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Loop Status</b>               | Indicates that a loopback issue is either found, not found, or unknown when loopback tracking is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>Detect LOC</b>                | Indicates that loss-of-continuity (LOC) detection is enabled or disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | All levels      |
| <b>LOC status</b>                | Indicates that a LOC issue is either found, not found, or unknown when Detect LOC is enabled. Status is unknown when LOC detection is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | All levels      |
| <b>OAM Receive Statistics</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                 |
| <b>Information</b>               | The total number of information PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>detail</b>   |
| <b>Event</b>                     | The total number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>   |
| <b>Variable request</b>          | The total number of variable request PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>detail</b>   |

Table 178: show oam ethernet link-fault-management Output Fields (*continued*)

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

Table 178: show oam ethernet link-fault-management Output Fields (*continued*)

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

Table 178: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                           | Field Description                                                                                                                                          | Level of Output |
|------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Threshold</b>                                     | The number of errored symbols in the period required for the event to be generated.                                                                        | <b>detail</b>   |
| <b>Errors in period</b>                              | The number of symbol errors in the period reported in the transmitted event PDU.                                                                           | <b>detail</b>   |
| <b>Total errors</b>                                  | The number of errored symbols reported in event TLVs that have been transmitted since the OAM sublayer was reset.                                          | <b>detail</b>   |
| <b>OAM Current Symbol Error Event Information</b>    |                                                                                                                                                            |                 |
| <b>Events</b>                                        | The number of symbol error TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.                      | <b>detail</b>   |
| <b>Window</b>                                        | The symbol error event window in the transmitted PDU.                                                                                                      | <b>detail</b>   |
| <b>Threshold</b>                                     | The number of errored symbols in the period required for the event to be generated.                                                                        | <b>detail</b>   |
| <b>Errors in period</b>                              | The total number of symbol errors in the period reported.                                                                                                  | <b>detail</b>   |
| <b>Total errors</b>                                  | The number of errored symbols reported in event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed. | <b>detail</b>   |
| <b>OAM Transmitted Frame Error Event Information</b> |                                                                                                                                                            |                 |
| <b>Events</b>                                        | The number of errored frame event TLVs that have been transmitted since the OAM sublayer was reset.                                                        | <b>detail</b>   |
| <b>Window</b>                                        | The duration of the window in terms of the number of 100 ms period intervals.                                                                              | <b>detail</b>   |
| <b>Threshold</b>                                     | The number of detected errored frames required for the event to be generated.                                                                              | <b>detail</b>   |
| <b>Errors in period</b>                              | The number of detected errored frames in the period.                                                                                                       | <b>detail</b>   |
| <b>Total errors</b>                                  | The number of errored frames that have been detected since the OAM sublayer was reset.                                                                     | <b>detail</b>   |



Table 178: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                       | Field Description                                                                                                                            | Level of Output |
|--------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>OAM Current Frame Error Event Information</b> |                                                                                                                                              |                 |
| <b>Events</b>                                    | The number of errored frame event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed. | <b>detail</b>   |
| <b>Window</b>                                    | The duration of the window in terms of the number of 100 ms period intervals.                                                                | <b>detail</b>   |
| <b>Threshold</b>                                 | The number of detected errored frames required for the event to be generated.                                                                | <b>detail</b>   |
| <b>Errors in period</b>                          | The number of errored frames in the period.                                                                                                  | <b>detail</b>   |
| <b>Total errors</b>                              | The number of errored frames detected regardless of whether the threshold for transmitting event TLVs has been crossed.                      | <b>detail</b>   |

## Sample Output

**show oam ethernet link-fault-management brief**

```
user@host> show oam ethernet link-fault-management brief
```

```
Interface: ge-3/1/3
  Status: Running, Discovery state: Send Any, ISSU
  Peer address: 00:90:69:72:2c:83
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50 Remote loopback status:
Disabled on local port, Enabled on peer port
  Remote entity information:
    Remote MUX action: discarding, Remote parser action: loopback
    Discovery mode: active, Unidirectional mode: unsupported
    Remote loopback mode: supported, Link events: supported
    Variable requests: unsupported, Remote in ISSU
```

**show oam ethernet link-fault-management brief (Loopback tracking)**

```
user@host> show oam ethernet link-fault-management
```

```

Interface: ge-3/1/3
  Status: Running, Discovery state: Active Send Local
  Peer address: 00:00:00:00:00:00
  Flags:0x8
  Loopback tracking: Enabled,      Loop Status: Found

```

### show oam ethernet link-fault-management detail

user@host> show oam ethernet link-fault-management detail

```

Interface: ge-6/1/0
  Status: Running, Discovery state: Send Any, ISSU
  Peer address: 00:90:69:0a:07:14
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
  OAM receive statistics:
    Information: 186365, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 0, Organization specific: 0
  OAM transmit statistics:
    Information: 186347, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 0, Organization specific: 0
  OAM received symbol error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame period error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM transmitted symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM transmitted frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  Remote entity information:
    Remote MUX action: forwarding, Remote parser action: forwarding
    Discovery mode: active, Unidirectional mode: unsupported

```

```
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported, Remote in ISSU
```

### show oam ethernet link-fault-management detail (backup Routing Engine)

user@host> show oam ethernet link-fault-management ge-0/2/0 detail

```
Interface: ge-0/2/0
  Status: Running, Discovery state: Send Any
  Transmit interval: 100ms, PDU threshold: 3 frames, Hold time: 300ms
  Peer address: ac:4b:c8:81:90:a4
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
  OAM receive statistics:
    Information: 0, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 0, Organization specific: 0
  OAM flags receive statistics:
    Critical event: 0, Dying gasp: 0, Link fault: 0
  OAM transmit statistics:
    Information: 0, Event: 0, Variable request: 0, Variable response: 0
    Loopback control: 786, Organization specific: 0
  OAM received symbol error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame period error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM received frame seconds error event information:
    Events: 0, Window: 0, Threshold: 0
    Errors in period: 0, Total errors: 0
  OAM transmitted symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current symbol error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM transmitted frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
  OAM current frame error event information:
    Events: 0, Window: 0, Threshold: 1
    Errors in period: 0, Total errors: 0
```

Loopback tracking: Enabled, Loop status: Not Found

Detect LOC: Enabled, LOC status: Not Found

Remote entity information:

Remote MUX action: forwarding, Remote parser action: forwarding

Discovery mode: active, Unidirectional mode: unsupported

Remote loopback mode: unsupported, Link events: supported

Variable requests: unsupported

Application profile statistics:

| Profile Name     | Invoked | Executed |
|------------------|---------|----------|
| LK_ADJ_LOSS100_1 | 1       | 1        |
| LK_ADJ_LOSS100_2 | 1       | 0        |
| LK_ADJ_LOSS100_3 | 1       | 0        |
| LK_ADJ_LOSS101_1 | 1       | 1        |
| LK_ADJ_LOSS101_2 | 1       | 0        |
| LK_ADJ_LOSS101_3 | 1       | 0        |
| LK_ADJ_LOSS106_1 | 0       | 0        |
| LK_ADJ_LOSS106_2 | 0       | 0        |
| LK_ADJ_LOSS106_3 | 0       | 0        |
| LK_ADJ_LOSS107_1 | 0       | 0        |
| LK_ADJ_LOSS107_2 | 0       | 0        |
| LK_ADJ_LOSS107_3 | 0       | 0        |

## show oam ethernet lmi

### Syntax

```
show oam ethernet lmi (interface <interface-name>)
```

### Release Information

Command introduced in Junos OS Release 9.5.

### Description

On routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet, and OAM Ethernet Local Management Interface (E-LMI) configuration, display the LMI information for the configured interfaces or optionally for a specified interface.

**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

**interface**—(Optional) Display LMI information for a specified interface.

**interface-name**—(Optional) Display Ethernet LMI information for the specified interface only.

### Required Privilege Level

View

### Output Fields

[Table 179 on page 1972](#) lists the output fields for the **show oam ethernet lmi** command. Output fields are listed in the approximate order in which they appear.

**Table 179: show oam ethernet lmi Output Fields**

| Field Name                        | Field Description                                                                                                          |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Physical Interface</b>         | Header for the EVC information showing the Ethernet virtual circuit (EVC) name, configuration, and active/inactive status. |
| <b>UNI Identifier</b>             | Name of the UNI.                                                                                                           |
| <b>EVC map type</b>               | EVC configuration.                                                                                                         |
| <b>Polling verification timer</b> | Polling verification timer status.                                                                                         |

Table 179: show oam ethernet lmi Output Fields (*continued*)

| Field Name                       | Field Description                                                                       |
|----------------------------------|-----------------------------------------------------------------------------------------|
| <b>E-LMI state</b>               | Operational status of the E-LMI configuration in the interfaces or specified interface. |
| <b>Priority/Untagged VLAN ID</b> | To be provided.                                                                         |
| <b>Default EVC</b>               | The EVC set as the default EVC.                                                         |
| <b>Associated EVCs</b>           | Heading for the list of configured EVCs.                                                |
| <b>EVC Identifier</b>            | EVC name.                                                                               |
| <b>Reference ID</b>              | To be provided.                                                                         |
| <b>Status</b>                    | Status active or not active.                                                            |
| <b>CE VLAN IDs</b>               | Customer edge VLAN ID numbers.                                                          |

## Sample Output

**show oam ethernet lmi interface**

user@host> **show oam ethernet lmi interface ge-1/1/1**

```
Physical interface: ge-1/1/1, Physical link is Up
UNI identifier: uni-cel, EVC map type: Bundling
Polling verification timer: Enabled, E-LMI state: Operational
Priority/Untagged VLAN ID: 20, Default EVC: evc1
Associated EVCs:
  EVC      Reference   Status      CE VLAN IDs
  Identifier ID
  evc1      1      Active (New)  1-2048
  evc2      2      Not Active   2049-4096
```

# show oam ethernet lmi statistics

## Syntax

```
show oam ethernet lmi statistics <interface interface-name>
```

## Release Information

Command introduced in Junos OS Release 9.5.

## Description

On MX Series routers with Gigabit Ethernet, Fast Ethernet, or aggregated Ethernet PICs, displays OAM Ethernet Local Management Interface (LMI) statistics.

## Options

**interface**—(Optional) Display LMI statistics for a specified interface.

**interface-name**—(Optional) Display Ethernet LMI information for the specified Ethernet interface only.

## Required Privilege Level

view

## List of Sample Output

[show oam ethernet lmi statistics on page 1975](#)

## Output Fields

[Table 180 on page 1974](#) lists the output fields for the **show oam ethernet lmi statistics** command. Output fields are listed in the approximate order in which they appear.

Table 180: show oam ethernet lmi statistics Output Fields

| Field Name            | Field Description                                   |
|-----------------------|-----------------------------------------------------|
| Physical interface    | Name of the interface for the displayed statistics. |
| Reliability errors    | Number of E-LMI reliability errors logged.          |
| Protocol errors       | Number of E-LMI protocol errorse.                   |
| Status check received | Number of E-LMI status check receive errors.        |
| Status check sent     | Number of E-LMI status check sent errors.           |
| Full status received  | Number of E-LMI full status receive errors.         |
| Full status sent      | Number of E-LMI full status sent errors.            |

Table 180: show oam ethernet lmi statistics Output Fields (*continued*)

| Field Name                            | Field Description                                  |
|---------------------------------------|----------------------------------------------------|
| <b>Full status continued received</b> | Number of E-LMI status continued received errors.  |
| <b>Full status continued sent</b>     | Number of E-LMI full status continued sent errors. |
| <b>Asynchronous status sent</b>       | Number of E-LMI asynchronous status sent errors.   |

## Sample Output

**show oam ethernet lmi statistics**

**user@host> show oam ethernet lmi statistics interface ge-1/1/1**

```
Physical interface: ge-1/1/1
  Reliability errors          4  Protocol errors
0
  Status check received      0  Status check sent
0
  Full status received       694  Full status sent          694

  Full status continued received  0  Full status continued sent
0
  Asynchronous status sent      0
```



## show pppoe interfaces

### Syntax

```
show pppoe interfaces
<brief | detail>
<pp0.logical>
```

### Release Information

Command introduced before Junos OS Release 7.4.

### 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

## RELATED DOCUMENTATION

*Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*

### List of Sample Output

[show pppoe interfaces on page 1978](#)

[show pppoe interfaces \(Status for the Specified Interface\) on page 1979](#)

[show pppoe interfaces brief on page 1979](#)

[show pppoe interfaces detail on page 1979](#)

[show pppoe interfaces \(PPPoE Subscriber Interface with ACI Interface Set\) on page 1980](#)

### Output Fields

[Table 181 on page 1977](#) 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 181: 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 none</b> |
| <b>State</b>                            | State of the logical interface: <b>up</b> or <b>down</b> .                                                                                                                                                                                                                                                                          | All levels                   |
| <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 none</b> |
| <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 none</b> |
| <b>Configured AC name</b>               | Configured access concentrator name.                                                                                                                                                                                                                                                                                                | <b>detail extensive none</b> |
| <b>Session AC name</b>                  | Name of the access concentrator.                                                                                                                                                                                                                                                                                                    | <b>detail extensive none</b> |
| <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 none</b> |
| <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 none</b> |
| <b>Underlying interface</b>             | Interface on which PPPoE is running.                                                                                                                                                                                                                                                                                                | All levels                   |
| <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 none</b> |

Table 181: show pppoe interfaces Output Fields (*continued*)

| Field Name               | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Level of Output              |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <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 none</b> |
| <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 none</b> |
| <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<br>interface | State      | Session<br>ID | Remote<br>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
```

# show pppoe service-name-tables

## Syntax

```
show pppoe service-name-tables
<table-name>
```

## Release Information

Command introduced in Junos OS Release 10.0.

## 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

## RELATED DOCUMENTATION

Verifying a PPPoE Configuration | 53

Verifying and Managing Dynamic PPPoE Configuration

## List of Sample Output

- [show pppoe service-name-tables on page 1983](#)
- [show pppoe service-name-tables \(For the Specified Table Name\) on page 1983](#)

## Output Fields

[Table 182 on page 1981](#) 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 182: show pppoe service-name-tables Output Fields

| Field Name         | Field Description                     | Level of Output |
|--------------------|---------------------------------------|-----------------|
| Service Name Table | Name of the PPPoE service name table. | none            |

Table 182: 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 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            |
| <b>Active Sessions</b>  | <p>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.</p>                                                                                                                                                                                                                                                                                                                                               | none            |
| <b>ACI</b>              | <p>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.</p>                                                                                                                                                            | none            |

Table 182: show pppoe service-name-tables Output Fields (*continued*)

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <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
```

## show pppoe sessions

### Syntax

```
show pppoe sessions  
<aci circuit-id-string>  
<ari remote-id-string>  
<service service-name>
```

### Release Information

Command introduced in Junos OS Release 10.2.

### 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

**none**—Display information for all active PPPoE sessions on the router.

**aci *circuit-id-string***—(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.

**ari *remote-id-string***—(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

## RELATED DOCUMENTATION

[Verifying a PPPoE Configuration | 53](#)

*Verifying and Managing Dynamic PPPoE Configuration*

### List of Sample Output

[show pppoe sessions \(For All Active Sessions\) on page 1986](#)

[show pppoe sessions \(For All Active Sessions Matching the Agent Circuit Identifier\) on page 1986](#)

### Output Fields

Table 183 on page 1986 lists the output fields for the **show pppoe sessions** command. Output fields are listed in the approximate order in which they appear.

**Table 183: 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 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 sessions (For All Active Sessions Matching the Agent Circuit Identifier)

```
user@host> show pppoe sessions aci "velorum-ge-2/0/3"
```

| 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 |

## show pppoe statistics

### Syntax

```
show pppoe statistics
<logical-interface-name>
```

### Release Information

Command introduced before Junos OS Release 7.4.

**logical-interface-name** option introduced in Junos OS Release 10.1.

### Description

Display statistics information about PPPoE interfaces.

### Options

**none**—Display PPPoE statistics for all interfaces.

**logical-interface-name**—(Optional) Name of a PPPoE underlying logical interface.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[show ppp address-pool](#)

[show pppoe underlying-interfaces](#) | [1990](#)

### List of Sample Output

[show pppoe statistics on page 1988](#)

[show pppoe statistics \(For the Specified Underlying Interface Only\) on page 1989](#)

### Output Fields

[Table 184 on page 1988](#) lists the output fields for the **show pppoe statistics** command. Output fields are listed in the approximate order in which they appear.

Table 184: 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        |

## show pppoe underlying-interfaces

### Syntax

```
show pppoe underlying-interfaces  
<brief | detail | extensive>  
<lockout>  
<logical-interface-name>
```

### Release Information

Command introduced in Junos OS Release 10.0.

**lockout** option added in Junos OS Release 11.4.

### 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

## 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*

### List of Sample Output

[show pppoe underlying-interfaces brief on page 1994](#)

[show pppoe underlying-interfaces detail on page 1995](#)

[show pppoe underlying-interfaces extensive on page 1995](#)

[show pppoe underlying-interfaces extensive \(PPPoE client in lockout condition\) on page 1996](#)

[show pppoe underlying-interfaces lockout on page 1997](#)

[show pppoe underlying-interfaces detail \(Autosensing Configured for ACI-based Dynamic VLANs\) on page 1997](#)

[show pppoe underlying-interfaces detail \(Autosensing Configured for ALI-based Dynamic VLANs\) on page 1998](#)

### Output Fields

Table 185 on page 1991 lists the output fields for the **show pppoe underlying-interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 185: 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>      |
| <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 none</b> |



Table 185: show pppoe underlying-interfaces Output Fields (continued)

| Field Name                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Level of Output              |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <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>      |
| <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 none</b> |
| <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 none</b> |
| <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. The uniqueness of the PPPoE client is determined by the client's MAC address.                                                                                                                                                                                                                                                                                                                                    | <b>detail extensive</b>      |

Table 185: 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 none</b> |
| <b>AC Name</b>                | Name of the access concentrator.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>detail extensive</b>      |
| <b>PacketType</b>             | Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: <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 185: 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,
```

# show pppoe version

## Syntax

```
show pppoe version
```

## Release Information

Command introduced before Junos OS Release 7.4.

## Description

( M120 routers and M320 routers only) Display version information about PPPoE.

## Options

This command has no options.

## Required Privilege Level

view

## List of Sample Output

[show pppoe version on page 2000](#)

## Output Fields

[Table 186 on page 1999](#) lists the output fields for the **show pppoe version** command. Output fields are listed in the approximate order in which they appear.

Table 186: show pppoe version Output Fields

| Field Name          | Field Description                                                                                                                                                                                                                                                                   |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| version <i>n</i>    | PPPoE version number and RFC. For example, <b>version 1, rfc 2516</b> .                                                                                                                                                                                                             |
| PPPoE protocol      | State of the PPPoE protocol: <b>enabled</b> or <b>disabled</b> .                                                                                                                                                                                                                    |
| Maximum Sessions    | Maximum active sessions supported per router. The default is 256 sessions.                                                                                                                                                                                                          |
| PADI resend timeout | 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.             |
| PADR resend timeout | 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. |



Table 186: show pppoe version Output Fields (*continued*)

| Field Name                | Field Description                                                                                                                           |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| Max resend timeout        | 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. |
| Max Configured AC timeout | 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
```

## show protection-group ethernet-ring aps

### Syntax

```
show protection-group ethernet-ring aps
```

### Release Information

Command introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 12.1 for EX Series switches.

Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

### Description

Display the status of the Automatic Protection Switching (APS) and Ring APS (RAPS) messages on an Ethernet ring.

### Options

This command has no options.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[show protection-group ethernet-ring data-channel | 2015](#)

[show protection-group ethernet-ring interface | 2020](#)

[show protection-group ethernet-ring node-state | 2025](#)

[show protection-group ethernet-ring statistics | 2031](#)

[show protection-group ethernet-ring vlan | 2038](#)

### List of Sample Output

[show protection-group ethernet-ring aps \(EX Switches\) on page 2003](#)

[show protection-group ethernet-ring aps \(Owner Node, Normal Operation on ACX and MX Routers\) on page 2003](#)

[show protection-group ethernet-ring aps detail \(Owner Node, Normal Operation on ACX and MX Routers\) on page 2003](#)

[show protection-group ethernet-ring aps \(MX RPL Owner Ring Node, Failure condition on non-RPL link of the ring\) on page 2003](#)

[show protection-group ethernet-ring aps \(MX Interconnection Ring Node, Failure condition in major ring on non-RPL link of the ring\) on page 2004](#)

[show protection-group ethernet-ring aps \(MX Series router\) on page 2004](#)

[show protection-group ethernet-ring aps detail \(MX Series router\) on page 2004](#)

[show protection-group ethernet-ring aps \(MX Interconnection Ring Node as RPL owner of major ring, rings in IDLE state\) on page 2005](#)

[show protection-group ethernet-ring aps detail \(EX2300 and EX3400 Switches\) on page 2005](#)

### Output Fields

[Table 187 on page 2002](#) lists the output fields for the **show protection-group ethernet-ring aps** command. Output fields are listed in the approximate order in which they appear.

**Table 187: show protection-group ethernet-ring aps Output Fields**

| Field Name                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ethernet Ring</b>                | Name configured for the Ethernet ring.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Request/State</b>                | <p>Status of the Ethernet ring RAPS messages.</p> <ul style="list-style-type: none"> <li>• <b>NR</b>—Indicates that there is no request for APS on the ring.</li> <li>• <b>SF</b>—Indicates that there is a signal failure on the ring.</li> <li>• <b>FS</b>—Indicates that there are active forced-switch requests in the ring.</li> <li>• <b>MS</b>—Indicates that there are active manual-switch requests in the ring.</li> </ul> <p><b>NOTE:</b> Both <b>FS</b> and <b>MS</b> values are valid only when G.8032v2 is supported.</p> |
| <b>Ring Protection Link Blocked</b> | Blocking on the ring protection link: <b>Yes</b> or <b>No</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>No Flush</b>                     | Indicates the value of the Do Not Flush (DNF) flag in the received RAPS PDU. If the value is Yes, then FDB flush is not triggered as part of processing of the received RAPS PDU.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Blocked Port Reference</b>       | This parameter is the reference to the blocked ring port. If the east ring port is blocked, the Blocked Port Reference (BPR) value is 0. If the west ring port is blocked, the BPR value is 1. If both ring ports are blocked, this parameter can take any value. If both east and west ports are blocked or not blocked, the value would be 0. This field is valid only when G.8032v2 is supported.                                                                                                                                    |
| <b>Blocked Port Reference</b>       | Reference of the ring port on which traffic is blocked.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Originator</b>                   | Indicates whether the node is the originator of the RAPS messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Remote Node ID</b>               | Identifier (in MAC address format) of the remote node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Sample Output

**show protection-group ethernet-ring aps (EX Switches)**

user@switch>show protection-group ethernet-ring aps

|           |               |          |             |            |                   |      |
|-----------|---------------|----------|-------------|------------|-------------------|------|
| Ring Name | Request/state | No Flush | RPL Blocked | Originator | Remote Node ID    | erp1 |
| NR        | No            | Yes      | No          |            | 00:1F:12:30:B8:81 |      |

## Sample Output

**show protection-group ethernet-ring aps (Owner Node, Normal Operation on ACX and MX Routers)**

user@host> show protection-group ethernet-ring aps

|                   |               |             |          |     |            |                |
|-------------------|---------------|-------------|----------|-----|------------|----------------|
| Ethernet Ring ID  | Request/state | RPL Blocked | No Flush | BPR | Originator | Remote Node ID |
| Erp_1             | NR            | Yes         | No       | 1   | No         |                |
| 00:00:00:02:00:01 |               |             |          |     |            |                |

## Sample Output

**show protection-group ethernet-ring aps detail (Owner Node, Normal Operation on ACX and MX Routers)**

user@host> show protection-group ethernet-ring aps detail

|                              |                     |
|------------------------------|---------------------|
| Ethernet-Ring name           | : Erp_1             |
| Request/State                | : NR                |
| Ring Protection Link blocked | : Yes               |
| No Flush Flag                | : No                |
| Blocked Port Reference       | : 1                 |
| Originator                   | : No                |
| Remote Node ID               | : 00:00:00:02:00:01 |

**show protection-group ethernet-ring aps (MX RPL Owner Ring Node, Failure condition on non-RPL link of the ring)**

user@host>show protection-group ethernet-ring aps

|               |               |             |          |
|---------------|---------------|-------------|----------|
| Ethernet Ring | Request/state | RPL Blocked | No Flush |
| pg101         | SF            | No          | No       |

|            |                   |
|------------|-------------------|
| Originator | Remote Node ID    |
| No         | 00:01:02:00:00:01 |

### show protection-group ethernet-ring aps (MX Interconnection Ring Node, Failure condition in major ring on non-RPL link of the ring)

user@host>show protection-group ethernet-ring aps

|               |               |             |          |     |
|---------------|---------------|-------------|----------|-----|
| Ethernet Ring | Request/state | RPL Blocked | No Flush | BPR |
| pg_major      | SF            | No          | No       | 0   |
| pg_subring    | NR            | Yes         | Yes      | 0   |

|            |                   |
|------------|-------------------|
| Originator | Remote Node ID    |
| No         | 00:01:00:00:00:01 |
| No         | 00:02:00:00:00:02 |

### show protection-group ethernet-ring aps (MX Series router)

user@host>show protection-group ethernet-ring aps

|                  |                   |             |          |     |            |                |
|------------------|-------------------|-------------|----------|-----|------------|----------------|
| Ethernet Ring    | Request/state     | RPL Blocked | No Flush | BPR | Originator | Remote Node ID |
| Inst_Vlans_1-15  | NR                | Yes         | Yes      | 1   | Yes        | NA             |
| Inst_Vlans_16-30 | NR                | Yes         | Yes      | 0   | No         |                |
|                  | 00:00:00:03:00:02 |             |          |     |            |                |

### show protection-group ethernet-ring aps detail (MX Series router)

user@host>show protection-group ethernet-ring aps

|                              |                   |
|------------------------------|-------------------|
| Ethernet-Ring name           | : Inst_Vlans_1-15 |
| Request/State                | : NR              |
| Ring Protection Link blocked | : Yes             |
| No Flush Flag                | : Yes             |
| Blocked Port Reference       | : 1               |
| Originator                   | : Yes             |
| Remote Node ID               | : NA              |

```

Ethernet-Ring name      : Inst_Vlans_16-30
Request/State           : NR
Ring Protection Link blocked : Yes
No Flush Flag           : Yes
Blocked Port Reference   : 0
Originator              : No
Remote Node ID          : 00:00:00:03:00:02

```

**show protection-group ethernet-ring aps (MX Interconnection Ring Node as RPL owner of major ring, rings in IDLE state)**

user@host>show protection-group ethernet-ring aps detail

```

Ethernet-Ring name      : pg_major
Request/State           : NR
Ring Protection Link blocked : Yes
No Flush Flag           : Yes
Blocked Port Reference   : 0
Originator              : Yes
Remote Node ID          : NA

Ethernet-Ring name      : pg_subring
Request/State           : NR
Ring Protection Link blocked : Yes
No Flush Flag           : Yes
Blocked Port Reference   : 0
Originator              : No
Remote Node ID          : 00:00:03:00:00:03

```

**show protection-group ethernet-ring aps detail (EX2300 and EX3400 Switches)**

user@switch>show protection-group ethernet-ring aps detail

```

Ethernet-Ring name      : pg1001
Request/State           : NR
Ring Protection Link blocked : Yes
No Flush Flag           : Yes
Blocked Port Reference   : 0
Originator              : Yes
Remote Node ID          : NA

```

# show protection-group ethernet-ring configuration

## Syntax

```
show protection-group ethernet-ring configuration
```

## Release Information

Command introduced in Junos OS Release 12.1 for EX Series switches.  
Command introduced in Junos OS Release 14.1 for MX Series routers.  
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

## Description

Display the configuration of Ethernet ring protection group on EX Switches and MX Series routers.

## Required Privilege Level

view

## RELATED DOCUMENTATION

|                                                                         |
|-------------------------------------------------------------------------|
| <a href="#">show protection-group ethernet-ring aps   2001</a>          |
| <a href="#">show protection-group ethernet-ring data-channel   2015</a> |
| <a href="#">show protection-group ethernet-ring interface   2020</a>    |
| <a href="#">show protection-group ethernet-ring node-state   2025</a>   |
| <a href="#">show protection-group ethernet-ring statistics   2031</a>   |
| <a href="#">show protection-group ethernet-ring vlan   2038</a>         |

## List of Sample Output

- [show protection-group ethernet-ring configuration \(EX Switch\) on page 2009](#)
- [show protection-group ethernet-ring configuration detail \(MX Series Router\) on page 2010](#)
- [show protection-group ethernet-ring configuration \(MX Series Router\) on page 2010](#)
- [show protection-group ethernet-ring configuration detail \(MX Series Router\) on page 2011](#)
- [show protection-group ethernet-ring configuration detail \(MX Series Router\) on page 2011](#)
- [show protection-group ethernet-ring configuration \(MX Series Router\) on page 2012](#)
- [show protection-group ethernet-ring configuration detail \(MX Series Router\) on page 2013](#)

## Output Fields

[Table 188 on page 2007](#) lists the output fields for the **show protection-group ethernet-ring configuration** command. Output fields are listed in the approximate order in which they appear.

Table 188: show protection-group ethernet-ring configuration Output Fields

| Output Fields                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>G8032 Compatability Version</b> | This is the compatibility version mode of ERP. This parameter always takes the value 1 in the case of G8032v1. This parameter is valid only for MX Series routers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>East Interface</b>              | One of the two switch interfaces that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>West Interface</b>              | One of the two interfaces in a switch that participates in a ring link. When Junos supports G8032v2, this interface is treated as interface 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Restore Interval</b>            | <p>Configured interval of wait time after a link is restored. When a link goes down, the RPL link is activated. When the down link becomes active again, the RPL owner receives a notification. The RPL owner waits for the restore interval before issuing a block on the RPL link. The configured restore interval can be 5 through 12 minutes for ER Pv1 and 1 through 12 minutes for ER Pv2. 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.</p> <p><b>NOTE:</b> Wait to Restore (WTR) configuration values on EX2300 and EX3400 switches must be 5-12 minutes.</p>               |
| <b>Wait to Block Interval</b>      | <p>Configured interval of wait time for link restoration when a manual command (manual switch or force switch) is cleared. On clearing the manual command, the RPL owner receives NR messages, which starts a timer with interval 'Wait to Block' to restore the RPL link after its expiration. This delay timer is set to be 5 seconds longer than the guard timer. The configured number can be from 5 seconds through 10 seconds. The parameter is valid only for G.8032v2.</p> <p><b>NOTE:</b> The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.</p> |



Table 188: show protection-group ethernet-ring configuration Output Fields (*continued*)

| Output Fields                      | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Guard Interval</b>              | Configured number of milliseconds (in 10 millisecond intervals, 10 milliseconds through 2000 milliseconds) that the node does not process any Ethernet ring protection 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. |
| <b>Hold off interval</b>           | This is the interval at which the link is held down even before declaring that the link is down. Because the parameter is not supported at present, its value is always considered 0. This parameter is valid only for MX Series routers.                                                                                                                                                                                                                                                          |
| <b>Node ID</b>                     | Node ID for the switch or router. If the node ID is not configured, it is assigned by default. For EX Series switches, the Node ID value cannot be configured, whereas for MX Series routers, it can be configured.                                                                                                                                                                                                                                                                                |
| <b>Ring ID</b>                     | In G8032v2, the ring ID can be within the range 1–239. All the nodes in a ring should have the same ring ID. In the case of G8032v1, the value of the ring ID is always 1. This parameter is valid only for MX Series routers.                                                                                                                                                                                                                                                                     |
| <b>Node Role</b>                   | Indicates whether the ring node is operating as a normal ring-node or RPL-owner or RPL-neighbor. For G8032v1 RPL-neighbor role is not supported. This parameter is valid only for MX Series routers.                                                                                                                                                                                                                                                                                               |
| <b>Revertive Mode of Operation</b> | This parameter indicates whether the ring is operating in revertive mode or nonrevertive mode. In nonrevertive mode of operation, when all links in the ring and Ethernet Ring Nodes have recovered and no external requests are active, the Ethernet Ring does not automatically revert. G8032v1 supports only revertive mode of operation. This parameter is valid only for MX Series routers.                                                                                                   |
| <b>RAPS Tx Dot1p priority</b>      | The RAPS Tx Dot1p priority is a parameter with which the RAPS is transmitted from the ring node. For G8032v1, the value of this parameter is always 0. For G8032v2, the value of this parameter can be within the range 0–7. This parameter is valid only for MX Series routers.                                                                                                                                                                                                                   |

Table 188: show protection-group ethernet-ring configuration Output Fields (*continued*)

| Output Fields                          | Field Description                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Node type</b>                       | Indicates whether ring node is a normal ring node having two ring-links or a open ring-node having only a single ring-link or a interconnection ring-node. An interconnection ring node can be connected to major ring in non virtual-channel mode or in virtual channel mode. Ring interconnection is not supported for G8032v1. This parameter is valid only for MX Series routers.                                    |
| <b>Major ring name</b>                 | If the node type is interconnection in the ring, this parameter takes the name of the major ring to which the sub-ring node is connected. This parameter is valid only for MX Series routers.                                                                                                                                                                                                                            |
| <b>Interconnection mode</b>            | Indicates the interconnection mode if the type of the node is interconnection. An interconnection ring node can be connected to major ring in non-virtual channel mode or in virtual channel mode. This parameter is valid only for MX Series routers.                                                                                                                                                                   |
| <b>Propagate Topology Change event</b> | When Propagate Topology Change event is set to 1, the change in the topology of sub-ring is propagated to the major ring, enabling the transmission of EVENT FLUSH RAPS PDU in the major ring. When the parameter is set to 0, the topology change in the sub-ring is not propagated to the major ring blocking EVENT FLUSH RAPS PDU transmission in the major ring. This parameter is valid only for MX Series routers. |
| <b>Control Vlan</b>                    | The VLAN that transfers ERP PDUs from one node to another.                                                                                                                                                                                                                                                                                                                                                               |
| <b>Physical Ring</b>                   | Physical ring if the east and west interfaces are nontrunk ports. For MX Series routers, the ring is termed a physical ring if no data channels are defined for the ring and the entire physical port forwarding is controlled by ERP.                                                                                                                                                                                   |
| <b>Data Channel VLAN(s)</b>            | Data VLANs for which forwarding behavior is controlled by the ring instance.                                                                                                                                                                                                                                                                                                                                             |

## Sample Output

show protection-group ethernet-ring configuration (EX Switch)

```
user@switch>show protection-group ethernet-ring configuration
```

```

Ethernet ring configuration parameters for protection group erp1
East Interface   : ge-0/0/3.0
West Interface   : ge-0/0/9.0
Restore Interval : 5 minutes
Guard Interval   : 500 ms
Node Id          : 00:1F:12:30:B8:81
Control Vlan     : 101
Physical Ring     : yes

```

### show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
```

```

Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version           : 2
East interface (interface 0)          : xe-2/3/0.1
West interface (interface 1)          : xe-2/2/1.1
Restore interval                      : 5 minutes
Wait to Block interval                : 5 seconds
Guard interval                       : 500 ms
Hold off interval                    : 0 ms
Node ID                              : 64:87:88:65:37:D0
Ring ID (1 ... 239)                  : 1
Node role (normal/rpl-owner/rpl-neighbour) : normal
Revertive mode of operation           : 1
RAPS Tx Dot1p priority (0 .. 7)      : 0
Node type (normal/open/interconnection) : Normal
Control Vlan                         : 100
Physical Ring                        : No
Data Channel Vlan(s)                 : 200,300

```

### show protection-group ethernet-ring configuration (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration
```

```

Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version           : 2
East interface (interface 0)          : xe-2/3/0.1
West interface (interface 1)          : xe-2/2/1.1
Restore interval                      : 5 minutes
Wait to Block interval                : 5 seconds
Guard interval                       : 500 ms
Hold off interval                    : 0 ms

```

```

Node ID                               : 64:87:88:65:37:D0
Ring ID (1 ... 239)                   : 1
Node role (normal/rpl-owner/rpl-neighbour) : rpl-neighbour
Node RPL end                           : east-port
Revertive mode of operation            : 1
RAPS Tx Dot1p priority (0 .. 7)       : 0
Node type (normal/open/interconnection) : Normal
Control Vlan                           : 100
Physical Ring                           : No
Data Channel Vlan(s)                   : 200,300

```

### show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
```

```

Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version           : 2
East interface (interface 0)          : xe-2/3/0.1
West interface (interface 1)          : xe-2/2/1.1
Restore interval                       : 5 minutes
Wait to Block interval                : 5 seconds
Guard interval                        : 500 ms
Hold off interval                     : 0 ms
Node ID                               : 64:87:88:65:37:D0
Ring ID (1 ... 239)                   : 1
Node role (normal/rpl-owner/rpl-neighbour) : rpl-owner
Node RPL end                           : east-port
Revertive mode of operation            : 1
RAPS Tx Dot1p priority (0 .. 7)       : 0
Node type (normal/open/interconnection) : Normal
Control Vlan                           : 100
Physical Ring                           : No
Data Channel Vlan(s)                   : 200,300

```

### show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
```

```

Ethernet Ring configuration information for protection group pg_101
G8032 Compatibility Version           : 2
East interface (interface 0)          : xe-2/3/0.1
West interface (interface 1)          : (no erp)
Restore interval                       : 5 minutes

```

```

Wait to Block interval          : 5 seconds
Guard interval                  : 500 ms
Hold off interval               : 0 ms
Node ID                         : 64:87:88:65:37:D0
Ring ID (1 ... 239)            : 1
Node role (normal/rpl-owner/rpl-neighbour) : rpl-owner
Node RPL end                    : east-port
Revertive mode of operation     : 1
RAPS Tx Dot1p priority (0 .. 7) : 0
Node type (normal/open/interconnection)   : Open
Control Vlan                    : 100
Physical Ring                   : No
Data Channel Vlan(s)            : 200,300

```

### show protection-group ethernet-ring configuration (MX Series Router)

user@switch>show protection-group ethernet-ring configuration

```

Ethernet Ring configuration information for protection group pg_major
G8032 Compatibility Version      : 2
East interface (interface 0)    : xe-2/3/0.1
West interface (interface 1)    : xe-2/2/1.1
Restore interval                 : 5 minutes
Wait to Block interval          : 5 seconds
Guard interval                  : 500 ms
Hold off interval               : 0 ms
Node ID                         : 64:87:88:65:37:D0
Ring ID (1 ... 239)            : 1
Node role (normal/rpl-owner/rpl-neighbour) : rpl-owner
Node RPL end                    : east-port
Revertive mode of operation     : 1
RAPS Tx Dot1p priority (0 .. 7) : 0
Node type (normal/open/interconnection)   : Normal
Control Vlan                    : 100
Physical Ring                   : No
Data Channel Vlan(s)            : 200,300

```

```

Ethernet Ring configuration information for protection group pg_subring
G8032 Compatibility Version      : 2
East interface (interface 0)    : ge-2/0/0.1
West interface (interface 1)    : (no erp)
Restore interval                 : 5 minutes
Wait to Block interval          : 5 seconds

```

```

Guard interval                : 500 ms
Hold off interval             : 0 ms
Node ID                       : 64:87:88:65:37:D0
Ring ID (1 ... 239)          : 2
Node role (normal/rpl-owner/rpl-neighbour) : normal
Revertive mode of operation   : 1
RAPS Tx Dot1p priority (0 .. 7) : 0
Node type (normal/open/interconnection)   : Non-VC-Interconnection
Major ring name               : pg_major
Interconnection mode (VC/Non-VC)          : Non-VC mode
Propagate Topology Change event           : 0
Control Vlan                     : 101
Physical Ring                    : No
Data Channel Vlan(s)             : 200,300

```

### show protection-group ethernet-ring configuration detail (MX Series Router)

```
user@switch>show protection-group ethernet-ring configuration detail
```

```

Ethernet Ring configuration information for protection group pg_major
G8032 Compatibility Version      : 2
East interface (interface 0)     : xe-2/3/0.1
West interface (interface 1)     : xe-2/2/1.1
Restore interval                 : 5 minutes
Wait to Block interval          : 5 seconds
Guard interval                  : 500 ms
Hold off interval               : 0 ms
Node ID                         : 64:87:88:65:37:D0
Ring ID (1 ... 239)            : 1
Node role (normal/rpl-owner/rpl-neighbour) : rpl-owner
Node RPL end                    : east-port
Revertive mode of operation     : 1
RAPS Tx Dot1p priority (0 .. 7) : 0
Node type (normal/open/interconnection)   : Normal
Control Vlan                    : 100
Physical Ring                   : No
Data Channel Vlan(s)           : 200,300

```

```

Ethernet Ring configuration information for protection group pg_subring
G8032 Compatibility Version      : 2
East interface (interface 0)     : ge-2/0/0.1
West interface (interface 1)     : (no erp)
Restore interval                 : 5 minutes

```

```
Wait to Block interval          : 5 seconds
Guard interval                  : 500 ms
Hold off interval               : 0 ms
Node ID                         : 64:87:88:65:37:D0
Ring ID (1 ... 239)            : 2
Node role (normal/rpl-owner/rpl-neighbour) : normal
Revertive mode of operation     : 1
RAPS Tx Dot1p priority (0 .. 7) : 0
Node type (normal/open/interconnection) : Non-VC-Interconnection
Major ring name                 : pg_major
Interconnection mode (VC/Non-VC) : Non-VC mode
Propagate Topology Change event : 0
Control Vlan                    : 101
Physical Ring                   : No
Data Channel Vlan(s)            : 200,300
```

# show protection-group ethernet-ring data-channel

## Syntax

```
show protection-group ethernet-ring data-channel
<brief | detail>
<group-name group-name>
```

## Release Information

Command introduced in Junos OS Release 10.2.  
Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

## Description

Display the configuration of Ethernet ring protection group on EX Switches and MX Series routers.

## Options

- brief | detail**—(Optional) Display the specified level of output.
- group-name**—(Optional) Protection group for which to display statistics. If you omit this optional field, all protection group statistics for configured groups will be displayed.

## Required Privilege Level

view

## RELATED DOCUMENTATION

|                                                                       |
|-----------------------------------------------------------------------|
| <a href="#">show protection-group ethernet-ring aps   2001</a>        |
| <a href="#">show protection-group ethernet-ring interface   2020</a>  |
| <a href="#">show protection-group ethernet-ring node-state   2025</a> |
| <a href="#">show protection-group ethernet-ring statistics   2031</a> |
| <a href="#">show protection-group ethernet-ring vlan   2038</a>       |

## List of Sample Output

- [show protection-group ethernet-ring data-channel on page 2016](#)
- [show protection-group ethernet-ring data-channel detail on page 2016](#)
- [show protection-group ethernet-ring data-channel detail \(EX2300 and EX3400 Switches\) on page 2017](#)

## Output Fields

[Table 189 on page 2016](#) lists the output fields for the **show protection-group ethernet-ring data-channel** command. Output fields are listed in the approximate order in which they appear.



Table 189: show protection-group ethernet-ring data-channel Output Fields

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>     | Name of the interface configured for the Ethernet ring.                                                                                                                                                                                                                                                                                                                                                                 |
| <b>STP index</b>     | The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior. |
| <b>Forward State</b> | Forwarding state on the Ethernet ring. <ul style="list-style-type: none"> <li>• <b>fowarding</b>—Indicates packets are being forwarded.</li> <li>• <b>discarding</b>—Indicates packets are being discarded.</li> </ul>                                                                                                                                                                                                  |

## Sample Output

### show protection-group ethernet-ring data-channel

```
user@host> show protection-group ethernet-ring data-channel
```

```
Ethernet ring data channel information for protection group pg301
```

| Interface | STP index | Forward State |
|-----------|-----------|---------------|
| xe-5/0/2  | 78        | forwarding    |
| xe-2/2/0  | 79        | discarding    |

```
Ethernet ring data channel parameters for protection group pg302
```

| Interface | STP index | Forward State |
|-----------|-----------|---------------|
| xe-5/0/2  | 80        | forwarding    |
| xe-2/2/0  | 81        | forwarding    |

### show protection-group ethernet-ring data-channel detail

```
user@host> show protection-group ethernet-ring data-channel detail
```

```
Ethernet ring data channel parameters for protection group pg301
```

```
Interface name           : xe-5/0/2
```

```

STP index           : 78
Forward State       : forwarding

Interface name      : xe-2/2/0
STP index           : 79
Forward State       : discarding

```

Ethernet ring data channel parameters for protection group pg302

```

Interface name      : xe-5/0/2
STP index           : 80
Forward State       : forwarding

```

```

Interface name      : xe-2/2/0
STP index           : 81
Forward State       : forwarding

```

### **show protection-group ethernet-ring data-channel detail (EX2300 and EX3400 Switches)**

```
user@switch>show protection-group ethernet-ring data-channel detail
```

Ethernet ring data channel parameters for protection group pg1001

```

Interface name      : ge-0/0/42
STP index           : 52
Forward State       : discarding

```

```

Interface name      : ge-0/0/38
STP index           : 53
Forward State       : forwarding

```

# show protection-group ethernet-ring flush-info

## Syntax

```
show protection-group ethernet-ring flush-info
```

## Release Information

Command introduced in Junos OS Release 14.2.

## Description

Display information about flush ports in an Ethernet ring.

## Options

This command has no options.

## Required Privilege Level

view

## RELATED DOCUMENTATION

- [show protection-group ethernet-ring data-channel | 2015](#)
- [show protection-group ethernet-ring aps | 2001](#)
- [show protection-group ethernet-ring node-state | 2025](#)
- [show protection-group ethernet-ring statistics | 2031](#)
- [show protection-group ethernet-ring vlan | 2038](#)

## List of Sample Output

- [show protection-group ethernet-ring flush-info \(ACX and MX Series Routers\) on page 2019](#)
- [show protection-group ethernet-ring flush-info detail \(ACX and MX Series Routers\) on page 2019](#)

## Output Fields

Table 190 on page 2018 lists the output fields for the **show protection-group ethernet-ring flush-info** command. Output fields are listed in the approximate order in which they appear.

Table 190: show protection-group ethernet-ring flush-info Output Fields

| Field Name       | Field Description                                                                                  |
|------------------|----------------------------------------------------------------------------------------------------|
| Interface        | Physical interface configured for the Ethernet ring. This can be an aggregated Ethernet link also. |
| Originating Node | Node from which RAPS protocol data units originates on the Ethernet Ring.                          |

Table 190: show protection-group ethernet-ring flush-info Output Fields (*continued*)

| Field Name                    | Field Description                                       |
|-------------------------------|---------------------------------------------------------|
| <b>Blocked Port Reference</b> | Reference of the ring port on which traffic is blocked. |

## Sample Output

### show protection-group ethernet-ring flush-info (ACX and MX Series Routers)

```
user@host> show protection-group ethernet-ring flush-info
```

```
Ethernet ring flush port information for protection group pg100
```

| Interface     | Originating Node  | Blocked Port Reference |
|---------------|-------------------|------------------------|
| xe-5/0/2.4001 | 00:00:00:00:00:00 | 0                      |
| xe-2/2/0.4001 | 00:00:00:00:00:00 | 0                      |

### show protection-group ethernet-ring flush-info detail (ACX and MX Series Routers)

```
user@host> show protection-group ethernet-ring flush-info detail
```

```
Ethernet ring flush port information for protection group pg100
```

|                        |                     |
|------------------------|---------------------|
| Interface name         | : xe-5/0/2.4001     |
| Originating Node       | : 00:00:00:00:00:00 |
| Blocked Port Reference | : 0                 |

|                        |                     |
|------------------------|---------------------|
| Interface name         | : xe-2/2/0.4001     |
| Originating Node       | : 00:00:00:00:00:00 |
| Blocked Port Reference | : 0                 |

## show protection-group ethernet-ring interface

### Syntax

```
show protection-group ethernet-ring interface
```

### Release Information

Command introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 12.3X54 for ACX Series routers.

Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

### Description

Displays the status of the Automatic Protection Switching (APS) interfaces on an Ethernet ring.

### Options

This command has no options.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[show protection-group ethernet-ring data-channel | 2015](#)

[show protection-group ethernet-ring aps | 2001](#)

[show protection-group ethernet-ring node-state | 2025](#)

[show protection-group ethernet-ring statistics | 2031](#)

[show protection-group ethernet-ring vlan | 2038](#)

### List of Sample Output

[show protection-group ethernet-ring interface \(EX Series Switch Owner Node\) on page 2021](#)

[show protection-group ethernet-ring interface \(Owner Node MX Series Router \) on page 2021](#)

[show protection-group ethernet-ring interface detail \(Owner Node MX Series Router \) on page 2022](#)

[show protection-group ethernet-ring interface \(EX Series Switch Ring Node\) on page 2022](#)

[show protection-group ethernet-ring interface detail \(ACX Series and MX Series\) on page 2023](#)

[show protection-group ethernet-ring interface detail \(EX2300 and EX3400 Switches\) on page 2023](#)

[show protection-group ethernet-ring interface detail \(EX2300 and EX3400 Switches\) on page 2024](#)

### Output Fields

[Table 191 on page 2021](#) lists the output fields for both the EX Series switch, and the ACX Series and MX Series router **show protection-group ethernet-ring interface** commands. Output fields are listed in the approximate order in which they appear.

Table 191: MX Series Routers show protection-group ethernet-ring interface Output Fields

| Field Name                                                           | Field Description                                                                                                                                        |
|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ethernet ring port parameters for protection group <i>group-name</i> | Output is organized by configured protection group.                                                                                                      |
| Interface                                                            | Physical interfaces configured for the Ethernet ring. This can be an aggregated Ethernet link also.                                                      |
| Control Channel                                                      | (MX Series router only) Logical unit configured on the physical interface.                                                                               |
| Direction                                                            | Direction of the traffic.                                                                                                                                |
| Forward State                                                        | State of the ring forwarding on the interface: <b>discarding</b> or <b>forwarding</b> .                                                                  |
| Ring Protection Link End                                             | Whether this interface is the end of the ring: <b>Yes</b> or <b>No</b> .                                                                                 |
| Signal Failure                                                       | Whether there a signal failure exists on the link: <b>Clear</b> or <b>Set</b> .                                                                          |
| Admin State                                                          | State of the interface: For EX switches, <b>ready</b> , <b>ifl ready</b> , or <b>waiting</b> . For MX routers, <b>IFF ready</b> or <b>IFF disabled</b> . |

## Sample Output

show protection-group ethernet-ring interface (EX Series Switch Owner Node)

```
user@host> show protection-group ethernet-ring interface
```

```
Ethernet ring port parameters for protection group pg101
```

| Interface  | Forward State | RPL End | Signal Failure | Admin State |
|------------|---------------|---------|----------------|-------------|
| ge-0/0/3.0 | discarding    | Yes     | Clear          | ready       |
| ge-0/0/9.0 | forwarding    | No      | Clear          | ready       |

show protection-group ethernet-ring interface (Owner Node MX Series Router )

```
user@host> show protection-group ethernet-ring interface
```

```
Ethernet ring port parameters for protection group pg101
```

| Interface | Control Channel | Direction | Forward State | RPL End | SF    | Admin State |
|-----------|-----------------|-----------|---------------|---------|-------|-------------|
| ge-1/2/0  | ge-1/2/0.100    | east      | forwarding    | No      | Clear | IFF ready   |
| ge-1/2/2  | ge-1/2/2.100    | west      | forwarding    | No      | Clear | IFF ready   |

### show protection-group ethernet-ring interface detail (Owner Node MX Series Router )

```
user@host> show protection-group ethernet-ring interface detail
```

```
Ethernet ring port parameters for protection group pg101
```

```
Interface name           : ge-1/2/0
Control channel name     : ge-1/2/0.100
Interface direction      : east
Ring Protection Link End : No
Signal Failure           : Clear
Forward State            : forwarding
Interface Admin State    : IFF ready
```

```
Interface name           : ge-1/2/2
Control channel name     : ge-1/2/2.100
Interface direction      : west
Ring Protection Link End : No
Signal Failure           : Clear
Forward State            : forwarding
Interface Admin State    : IFF ready
```

### show protection-group ethernet-ring interface (EX Series Switch Ring Node)

```
user@host> show protection-group ethernet-ring interface
```

```
Ethernet ring port parameters for protection group pg102
```

```
Ethernet ring port parameters for protection group pg101
```

| Interface  | Forward State | RPL End | Signal Failure | Admin State |
|------------|---------------|---------|----------------|-------------|
| ge-0/0/3.0 | discarding    | Yes     | Clear          | ready       |
| ge-0/0/9.0 | forwarding    | No      | Clear          | ready       |

**show protection-group ethernet-ring interface detail (ACX Series and MX Series)**

```
user@host> show protection-group ethernet-ring interface detail
```

```
Ethernet ring port parameters for protection group Erp_1
```

```
Interface name           : xe-0/0/0
Control channel name     : xe-0/0/0.1
Interface direction      : east
Ring Protection Link End : No
Signal Failure           : Clear
Forward State            : forwarding
Interface Admin State    : IFF ready
```

```
Interface name           : et-0/0/48
Control channel name     : et-0/0/48.1
Interface direction      : west
Ring Protection Link End : No
Signal Failure           : Clear
Forward State            : forwarding
Interface Admin State    : IFF ready
```

**show protection-group ethernet-ring interface detail (EX2300 and EX3400 Switches)**

```
user@switch>show protection-group ethernet-ring interface detail
```

```
Ethernet ring port parameters for protection group pg1001
```

```
Interface name           : ge-0/0/14
Control channel name     : ge-0/0/14.0
Interface direction      : east
Ring Protection Link End : No
Signal Failure           : Clear
Forward State            : forwarding
Interface Admin State    : IFF ready
```

```
Interface name           : ge-0/0/18
Control channel name     : ge-0/0/18.0
Interface direction      : west
Ring Protection Link End : No
Signal Failure           : Clear
Forward State            : forwarding
Interface Admin State    : IFF ready
```



**show protection-group ethernet-ring interface detail (EX2300 and EX3400 Switches)**

user@switch>show protection-group ethernet-ring interface detail

Ethernet ring port parameters for protection group pg1001

|                          |               |
|--------------------------|---------------|
| Interface name           | : ge-0/0/42   |
| Control channel name     | : ge-0/0/42.0 |
| Interface direction      | : east        |
| Ring Protection Link End | : Yes         |
| Signal Failure           | : Clear       |
| Forward State            | : discarding  |
| Interface Admin State    | : IFF ready   |

|                          |               |
|--------------------------|---------------|
| Interface name           | : ge-0/0/38   |
| Control channel name     | : ge-0/0/38.0 |
| Interface direction      | : west        |
| Ring Protection Link End | : No          |
| Signal Failure           | : Clear       |
| Forward State            | : forwarding  |
| Interface Admin State    | : IFF ready   |

## show protection-group ethernet-ring node-state

### Syntax

```
show protection-group ethernet-ring node-state
```

### Release Information

Command introduced in Junos OS Release 9.4 for MX Series routers.

Command introduced in Junos OS Release 12.1 for EX Series switches.

Command introduced in Junos OS Release 12.3X54 for ACX Series routers.

Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

### Description

Display the status of the Automatic Protection Switching (APS) nodes on an Ethernet ring.

### Options

This command has no options.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[show protection-group ethernet-ring data-channel | 2015](#)

[show protection-group ethernet-ring aps | 2001](#)

[show protection-group ethernet-ring interface | 2020](#)

[show protection-group ethernet-ring statistics | 2031](#)

[show protection-group ethernet-ring vlan | 2038](#)

### List of Sample Output

[show protection-group ethernet-ring node-state \(MX Series Router - RPL Owner Node, Normal Operation\) on page 2027](#)

[show protection-group ethernet-ring node-state \(MX Series Router - Normal Ring Node, Normal Operation\) on page 2028](#)

[show protection-group ethernet-ring node-state \(MX Series Router - RPL Owner Node, Remote Failure Condition\) on page 2028](#)

[show protection-group ethernet-ring node-state detail \(ACX Series and MX Series Router\) on page 2028](#)

[show protection-group ethernet-ring node-state detail \(MX Series Router - RPL Owner Node, Normal Operation\) on page 2028](#)

[show protection-group ethernet-ring node-state detail \(MX Series Router with WTR Timer\) on page 2029](#)

[show protection-group ethernet-ring node-state detail \(MX Series Router with WTB Timer\) on page 2029](#)

[show protection-group ethernet-ring node-state detail \(EX2300 and EX3400 Switches\) on page 2030](#)

## Output Fields

[Table 192 on page 2026](#) lists the output fields for the **show protection-group ethernet-ring node-state** command. Output fields are listed in the approximate order in which they appear.

**Table 192: show protection-group ethernet-ring node-state Output Fields**

| Field Name              | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ring Name/Ethernet Ring | Name configured for the Ethernet ring.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| APS State               | <p>State of the Ethernet ring APS.</p> <ul style="list-style-type: none"> <li>• <b>idle</b>—Indicates that the ring is working in normal condition and there is no active or pending protection-switching request in the ring. When the ring is in idle state, it is blocked at the RPL link.</li> <li>• <b>protected</b>—Indicates that there is a protection switch on the ring because of a signal failure condition on the ring link.</li> <li>• <b>MS</b>—Indicates that the manual switch command is active in the ring.</li> <li>• <b>FS</b>—Indicates that the forced switch command is active in the ring.</li> <li>• <b>pending</b>—Indicates that the ring is in pending state.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Event                   | <p>Events on the ring.</p> <ul style="list-style-type: none"> <li>• <b>NR-RB</b>—Indicates that there is no APS request and the ring link is blocked on the ring owner node.</li> <li>• <b>NR</b>—Indicates that there is no APS request pending in the ring.</li> <li>• <b>local SF</b>—Indicates that there is signal failure on one or both of the ring links of the node.</li> <li>• <b>remote SF</b>—Indicates that there is signal failure on one or more ring links of any other node of the ring.</li> <li>• <b>local FS</b>—Indicates that there is a forced switched command active on one or both of the ring links of the node.</li> <li>• <b>remote FS</b>—Indicates that there is a forced switch command active on one or more ring links of any other node of the ring.</li> <li>• <b>local MS</b>—Indicates that there is a manual switch command active on one of the ring links of the node.</li> <li>• <b>remote MS</b>—Indicates that there is a manual switch command active on one or more ring links of any other node of the ring.</li> <li>• <b>WTR running</b>—Indicates that the wait to restore timer is running on the RPL owner.</li> <li>• <b>WTB running</b>—Indicates that the wait to block timer is running on the RPL owner.</li> </ul> |

Table 192: show protection-group ethernet-ring node-state Output Fields (continued)

| Field Name                             | Field Description                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RPL Owner / Ring Protection Link Owner | Whether this node is the ring owner: <b>Yes</b> or <b>No</b> .                                                                                                                                                                                                                                                                                                                                              |
| WTR Timer / Restore Timer              | Restoration timer: <b>running</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                      |
| WTB Timer / Wait to block timer        | Wait to block timer: <b>running</b> or <b>disabled</b> .<br><br><b>NOTE:</b> The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect. |
| Wait to block timer (WTB Timer)        | Wait to block interval.<br><br><b>NOTE:</b> The Wait To Block Timer (WTB) is always disabled on EX2300 and EX3400 switches because it is not supported in ERPSv1. Any configuration you make to the WTB setting has no effect. The output from the CLI command 'show protection-group ethernet-ring node-state detail' lists a WTB setting but that setting has no effect.                                  |
| Guard Timer                            | Guard timer: <b>running</b> or <b>disabled</b> .                                                                                                                                                                                                                                                                                                                                                            |
| Op State / Operational State           | State of the node: <b>Operational</b> or <b>any internal wait state..</b>                                                                                                                                                                                                                                                                                                                                   |

## Sample Output

show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Normal Operation)

```
user@host> show protection-group ethernet-ring node-state
```

| Ethernet ring   | APS State | Event | RPL Owner | WTR Timer | WTB Timer | Guard Timer |
|-----------------|-----------|-------|-----------|-----------|-----------|-------------|
| Operation state |           |       |           |           |           |             |
| pg101           | idle      | NR-RB | Yes       | disabled  | disabled  | disabled    |
| operational     |           |       |           |           |           |             |
| pg102           | idle      | NR-RB | No        | disabled  | disabled  | disabled    |
| operational     |           |       |           |           |           |             |

**show protection-group ethernet-ring node-state (MX Series Router - Normal Ring Node, Normal Operation)**

```
user@host> show protection-group ethernet-ring node-state
```

|               |           |             |                 |
|---------------|-----------|-------------|-----------------|
| Ethernet ring | APS State | Event       | RPL Owner       |
| pg102         | idle      | NR-RB       | No              |
| WTR Timer     | WTB Timer | Guard Timer | Operation state |
| disabled      | disabled  | disabled    | operational     |

**show protection-group ethernet-ring node-state (MX Series Router - RPL Owner Node, Remote Failure Condition)**

```
user@host> show protection-group ethernet-ring node-state
```

|               |           |             |                 |
|---------------|-----------|-------------|-----------------|
| Ethernet ring | APS State | Event       | RPL Owner       |
| pg101         | protected | remote SF   | Yes             |
| WTR Timer     | WTB Timer | Guard Timer | Operation state |
| disabled      | disabled  | disabled    | operational     |

**show protection-group ethernet-ring node-state detail (ACX Series and MX Series Router)**

```
user@host> show protection-group ethernet-ring node-state detail
```

|                            |               |
|----------------------------|---------------|
| Ethernet-Ring name         | : Erp_1       |
| APS State                  | : idle        |
| Event                      | : NR-RB       |
| Ring Protection Link Owner | : No          |
| Wait to Restore Timer      | : disabled    |
| Wait to Block Timer        | : disabled    |
| Guard Timer                | : disabled    |
| Operation state            | : operational |

**show protection-group ethernet-ring node-state detail (MX Series Router - RPL Owner Node, Normal Operation)**

```
user@host> show protection-group ethernet-ring node-state detail
```

|                            |            |
|----------------------------|------------|
| Ethernet-Ring name         | : pg101    |
| APS State                  | : idle     |
| Event                      | : NR-RB    |
| Ring Protection Link Owner | : Yes      |
| Wait to Restore Timer      | : disabled |

```

Wait to Block Timer      : disabled
Guard Timer             : disabled
Operation state          : operational

```

```

Ethernet-Ring name      : pg102
APS State                : idle
Event                   : NR-RB
Ring Protection Link Owner : No
Wait to Restore Timer    : disabled
Wait to Block Timer      : disabled
Guard Timer             : disabled
Operation state          : operational

```

### show protection-group ethernet-ring node-state detail (MX Series Router with WTR Timer)

```
user@host> show protection-group ethernet-ring node-state detail
```

```

Ethernet-Ring name      : pg_major
APS State                : pending
Event                   : WTR running
Ring Protection Link Owner : Yes
Wait to Restore Timer    : running (time to expire: 269 sec)
Wait to Block Timer      : disabled
Guard Timer             : disabled
Operation state          : operational

```

```

Ethernet-Ring name      : pg_subring
APS State                : pending
Event                   : NR
Ring Protection Link Owner : No
Wait to Restore Timer    : disabled
Wait to Block Timer      : disabled
Guard Timer             : disabled
Operation state          : operational

```

### show protection-group ethernet-ring node-state detail (MX Series Router with WTB Timer)

```
user@host> show protection-group ethernet-ring node-state detail
```

```

Ethernet-Ring name      : Pg-2
APS State                : pending
Event                   : WTB running
Ring Protection Link Owner : Yes

```

```

Wait to Restore Timer      : disabled
Wait to Block Timer       : running (time to expire: 2 sec)
Guard Timer               : disabled
Operation state           : operational

```

### **show protection-group ethernet-ring node-state detail (EX2300 and EX3400 Switches)**

user@switch>show protection-group ethernet-ring node-state detail

```

Ethernet-Ring name        : pg1001
APS State                 : idle
Event                    : NR-RB
Ring Protection Link Owner : Yes
Wait to Restore Timer     : disabled
Wait to Block Timer       : disabled  <-field not supported. Always disabled.
Guard Timer              : disabled
Operation state           : operational

```

## show protection-group ethernet-ring statistics

### Syntax

```
show protection-group ethernet-ring statistics group-name group-name
```

<brief | detail>

### Release Information

Command introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 12.1 for EX Series switches.

Command introduced in Junos OS Release 12.3X54 for ACX Series routers.

### Description

Display statistics regarding Automatic Protection Switching (APS) protection groups on an Ethernet ring.

### Options

**group-name**—Display statistics for the protection group. If you omit this option, protection group statistics for all configured groups are displayed.

**brief**—Display brief statistics for the protection group.

**detail**—Display detailed statistics for the protection group.

### Required Privilege Level

view

## RELATED DOCUMENTATION

[show protection-group ethernet-ring data-channel | 2015](#)

[show protection-group ethernet-ring aps | 2001](#)

[show protection-group ethernet-ring node-state | 2025](#)

[show protection-group ethernet-ring interface | 2020](#)

[show protection-group ethernet-ring vlan | 2038](#)

### List of Sample Output

[show protection-group ethernet-ring statistics \(EX Series Switch\) on page 2034](#)

[show protection-group ethernet-ring statistics \(MX Series Router\) on page 2034](#)

[show protection-group ethernet-ring statistics detail \(Specific Group\)\(MX Series Router\) on page 2034](#)

[show protection-group ethernet-ring statistics \(Owner Node, Failure Condition on ACX and MX Router\) on page 2035](#)



[show protection-group ethernet-ring statistics \(Ring Node, Failure Condition on ACX and MX Router\) on page 2036](#)

[show protection-group ethernet-ring statistics detail \(EX2300 and EX3400 Switches\) on page 2036](#)

[show protection-group ethernet-ring statistics detail \(EX2300 and EX3400 Switches\) on page 2037](#)

### Output Fields

Table 193 on page 2032 lists the output fields for the **show protection-group ethernet-ring statistics** command.

Table 193: show protection-group ethernet-ring statistics Output Fields

| Field Name                             | Field Description                                                                                                                                               |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Ethernet Ring Statistics for PG</b> | Name of the protection group for which statistics are displayed.                                                                                                |
| <b>RAPS event sent</b>                 | Number of times Ring Automatic Protection Switching (RAPS) message transmission event occurred locally. This field is applicable only to MX Series routers.     |
| <b>RAPS event received</b>             | Number of RAPS messages received and processed by ERP state-machine and which resulted in state transition. This field is applicable only to MX Series routers. |
| <b>Local SF</b>                        | Number of times a signal failure has occurred locally.                                                                                                          |
| <b>Remote SF</b>                       | Number of times a signal failure has occurred anywhere else on the ring.                                                                                        |
| <b>NR event</b>                        | Number of times a No Request event has occurred on the ring. This field is applicable only to EX Series switches.                                               |
| <b>NR event sent</b>                   | Number of times a No Request event has occurred locally. This field is applicable only to MX Series routers.                                                    |
| <b>NR event received</b>               | Number of times a No Request event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.                                  |
| <b>NR-RB event</b>                     | Number of times a No Request, Ring Blocked event has occurred on the ring. This field is applicable only to EX Series switches.                                 |
| <b>NR-RB event sent</b>                | Number of times a No Request, Ring Blocked event has occurred locally. This field is applicable only to MX Series routers.                                      |
| <b>NR-RB event received</b>            | Number of times a No Request, Ring Blocked event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.                    |

Table 193: show protection-group ethernet-ring statistics Output Fields (*continued*)

| Field Name                      | Field Description                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flush event sent</b>         | Number of times flush-event RAPS message transmission event occurred locally. This field is applicable only to MX Series routers.                    |
| <b>Flush event received</b>     | Number of flush-event RAPS messages received and processed by the ring instance control process. This field is applicable only to MX Series routers. |
| <b>Local FS event sent</b>      | Number of times a forced switch event has occurred locally. This field is applicable only to MX Series routers.                                      |
| <b>Remote FS event received</b> | Number of times a forced switch event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.                    |
| <b>Local MS event sent</b>      | Number of times a manual switch event has occurred locally. This field is applicable only to MX Series routers.                                      |
| <b>Remote MS event received</b> | Number of times a manual switch event has occurred anywhere else on the ring. This field is applicable only to MX Series routers.                    |

Table 194 on page 2033 lists the output fields for the **show protection-group ethernet-ring statistics** command when the **detail** option is used. These fields are valid only for MX Series routers.

Table 194: show protection-group ethernet-ring statistics detail Output Fields (for MX Series Routers)

| Field Name                                      | Field Description                                                                                                                |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Total number of FDB flush</b>                | Number of times forwarding database (FDB) flush has happened for the ring instance.                                              |
| <b>Flush-logic triggered flush</b>              | Number of times FDB flush has happened because of flush-logic based on node ID and Blocked Port Reference (BPR).                 |
| <b>Remote RAPS PDU received</b>                 | Number of valid RAPS PDU messages received. This counter counts only RAPS messages generated by other devices on the ring.       |
| <b>Remote RAPS dropped due to guard-timer</b>   | Number of RAPS messages dropped by the device because the guard timer is running.                                                |
| <b>Invalid remote RAPS PDU dropped</b>          | Number of RAPS messages dropped by the device because the messages are invalid.                                                  |
| <b>RAPS dropped due to miscellaneous errors</b> | Number of RAPS messages dropped because of any other reason. For example, messages dropped because of unsupported functionality. |

**Table 194: show protection-group ethernet-ring statistics detail Output Fields (for MX Series Routers) (continued)**

| Field Name                      | Field Description                                            |
|---------------------------------|--------------------------------------------------------------|
| Local received RAPS PDU dropped | Number of self-generated RAPS messages received and dropped. |

## Sample Output

### show protection-group ethernet-ring statistics (EX Series Switch)

```
user@switch> show protection-group ethernet-ring statistics
```

```
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1      2      1      2      3
```

### show protection-group ethernet-ring statistics (MX Series Router)

```
user@host> show protection-group ethernet-ring statistics
```

```
Ethernet Ring statistics for PG Pg-1
RAPS event sent           : 1
RAPS event received       : 1152
Local SF happened:        : 0
Remote SF happened:       : 428
NR event sent:           : 1
NR event received:       : 133
NR-RB event sent:        : 0
NR-RB event received:    : 591
Flush event sent         : 0
Flush event received:    : 0
Local FS event sent:     : 0
Remote FS event received: : 0
Local MS event sent:     : 0
Remote MS event received: : 0
```

### show protection-group ethernet-ring statistics detail (Specific Group)(MX Series Router)

```
user@host> show protection-group ethernet-ring statistics detail
```

```

Ethernet Ring statistics for PG Pg-1
RAPS event sent                : 1
RAPS event received            : 0
Local SF happened              : 0
Remote SF happened             : 0
NR event sent                  : 1
NR event received              : 0
NR-RB event sent               : 0
NR-RB event received           : 0
Flush event sent               : 0
Flush event received           : 0
Local FS event sent            : 0
Remote FS event received       : 0
Local MS event sent            : 0
Remote MS event received       : 0
Total number of FDB flush     : 0
Flush-logic triggered flush    : 0
Remote raps PDU received       : 0
Remote raps dropped due to guard-timer : 0
Invalid remote raps PDU dropped : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped : 0

```

**show protection-group ethernet-ring statistics (Owner Node, Failure Condition on ACX and MX Router)**

user@host> **show protection-group ethernet-ring statistics group-name pg101**

```

Ethernet Ring statistics for PG pg101
RAPS sent                      : 1
RAPS received                  : 0
Local SF happened:             : 0
Remote SF happened:            : 0
NR event happened:              : 0
NR-RB event happened:           : 1
NR event sent:                 : 0
NR event received:             : 0
NR-RB event sent:              : 1
NR-RB event received:          : 0
Flush event sent               : 0
Flush event received:          : 0
Local FS event sent:           : 0
Remote FS event received:      : 0
Local MS event sent:           : 0
Remote MS event received:      : 0

```

**show protection-group ethernet-ring statistics (Ring Node, Failure Condition on ACX and MX Router)**

```
user@host> show protection-group ethernet-ring statistics group-name pg102
```

```
Ethernet Ring statistics for PG pg102
RAPS sent                : 1
RAPS received            : 0
Local SF happened:       : 0
Remote SF happened:      : 0
NR event happened:       : 0
NR-RB event happened:    : 1
NR event sent:          : 0
NR event received:      : 0
NR-RB event sent:       : 1
NR-RB event received:   : 0
Flush event sent        : 0
Flush event received:   : 0
Local FS event sent:    : 0
Remote FS event received: : 0
Local MS event sent:    : 0
Remote MS event received: : 0
```

**show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches)**

```
user@switch>show protection-group ethernet-ring statistics detail
```

```
Ethernet Ring statistics for PG pg1001
RAPS event sent          : 1
RAPS event received      : 1
Local SF happened        : 0
Remote SF happened       : 0
NR event sent           : 1
NR event received        : 0
NR-RB event sent         : 0
NR-RB event received     : 1
Flush event sent         : 0
Flush event received     : 0
Local FS event sent      : 0
Remote FS event received : 0
Local MS event sent      : 0
Remote MS event received : 0
Total number of FDB flush : 0
Flush-logic triggered flush : 0
Remote raps PDU received : 145
Remote raps dropped due to guard-timer : 0
```

```

Invalid remote raps PDU dropped      : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped      : 0

```

### **show protection-group ethernet-ring statistics detail (EX2300 and EX3400 Switches)**

user@switch>**show protection-group ethernet-ring statistics detail**

```

Ethernet Ring statistics for PG pg1001
RAPS event sent                : 2
RAPS event received            : 0
Local SF happened              : 0
Remote SF happened             : 0
NR event sent                  : 1
NR event received              : 0
NR-RB event sent               : 1
NR-RB event received           : 0
Flush event sent               : 0
Flush event received           : 0
Total number of FDB flush      : 0
Remote raps PDU received       : 211
Remote raps dropped due to guard-timer : 0
Invalid remote raps PDU dropped : 0
Raps dropped due to miscellaneous errors : 0
Local received raps PDU dropped : 91

```

## show protection-group ethernet-ring vlan

### Syntax

```
show protection-group ethernet-ring vlan  
<brief | detail>  
<group-name group-name>
```

### Release Information

Command introduced in Junos OS Release 10.2.

Command introduced in Junos OS Release 18.1 for EX2300 and EX3400 switches.

### Description

On MX Series routers, display all data channel logical interfaces and the VLAN IDs controlled by a ring instance data channel.

### Options

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

**group-name**—(Optional) Protection group for which to display details such as data channel interfaces, vlan, and bridge-domain. If you omit this optional field, details for all configured protection groups will be displayed.

### Required Privilege Level

view

### RELATED DOCUMENTATION

[show protection-group ethernet-ring aps | 2001](#)

[show protection-group ethernet-ring data-channel | 2015](#)

[show protection-group ethernet-ring interface | 2020](#)

[show protection-group ethernet-ring node-state | 2025](#)

[show protection-group ethernet-ring statistics | 2031](#)

### List of Sample Output

[show protection-group ethernet-ring vlan on page 2039](#)

[show protection-group ethernet-ring vlan brief on page 2040](#)

[show protection-group ethernet-ring vlan detail on page 2041](#)

[show protection-group ethernet-ring vlan group-name vkm01 on page 2042](#)

[show protection-group ethernet-ring vlan detail \(EX2300 and EX3400 Switches\) on page 2042](#)

## Output Fields

Table 195 on page 2039 lists the output fields for the **show protection-group ethernet-ring vlan** command. Output fields are listed in the approximate order in which they appear.

Table 195: show protection-group ethernet-ring vlan Output Fields

| Field Name           | Field Description                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Interface</b>     | Name of the interface configured for the Ethernet protection ring.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Vlan</b>          | Name of the VLAN associated with the interface configured for the Ethernet protection ring.                                                                                                                                                                                                                                                                                                                             |
| <b>STP index</b>     | The Spanning Tree Protocol (STP) index number used by each interface in an Ethernet ring. The STP index controls the forwarding behavior for a set of VLANs on a data channel on an Ethernet ring port. For multiple Ethernet ring instances on an physical ring port, there are multiple STP index numbers. Different ring instances will have different STP index numbers and may have different forwarding behavior. |
| <b>Bridge Domain</b> | Name of the bridge domain that is associated with the VLAN configured for the Ethernet protection ring.                                                                                                                                                                                                                                                                                                                 |

## Sample Output

**show protection-group ethernet-ring vlan**

```
user@host> show protection-group ethernet-ring vlan
```

```
Ethernet ring IFBD parameters for protection group vkm01
```

| Interface | Vlan | STP Index | Bridge Domain      |
|-----------|------|-----------|--------------------|
| xe-5/0/2  | 1    | 78        | default-switch/bd1 |
| xe-2/2/0  | 1    | 79        | default-switch/bd1 |
| xe-5/0/2  | 2    | 78        | default-switch/bd2 |
| xe-2/2/0  | 2    | 79        | default-switch/bd2 |
| xe-5/0/2  | 3    | 78        | default-switch/bd3 |
| xe-2/2/0  | 3    | 79        | default-switch/bd3 |
| xe-5/0/2  | 4    | 78        | default-switch/bd4 |
| xe-2/2/0  | 4    | 79        | default-switch/bd4 |
| xe-5/0/2  | 5    | 78        | default-switch/bd5 |
| xe-2/2/0  | 5    | 79        | default-switch/bd5 |
| xe-5/0/2  | 6    | 78        | default-switch/bd6 |



|          |    |    |                     |
|----------|----|----|---------------------|
| xe-2/2/0 | 6  | 79 | default-switch/bd6  |
| xe-5/0/2 | 7  | 78 | default-switch/bd7  |
| xe-2/2/0 | 7  | 79 | default-switch/bd7  |
| xe-5/0/2 | 8  | 78 | default-switch/bd8  |
| xe-2/2/0 | 8  | 79 | default-switch/bd8  |
| xe-5/0/2 | 9  | 78 | default-switch/bd9  |
| xe-2/2/0 | 9  | 79 | default-switch/bd9  |
| xe-5/0/2 | 10 | 78 | default-switch/bd10 |
| xe-2/2/0 | 10 | 79 | default-switch/bd10 |
| xe-5/0/2 | 11 | 78 | default-switch/bd11 |
| xe-2/2/0 | 11 | 79 | default-switch/bd11 |
| xe-5/0/2 | 12 | 78 | default-switch/bd12 |
| xe-2/2/0 | 12 | 79 | default-switch/bd12 |
| xe-5/0/2 | 13 | 78 | default-switch/bd13 |
| xe-2/2/0 | 13 | 79 | default-switch/bd13 |
| xe-5/0/2 | 14 | 78 | default-switch/bd14 |
| xe-2/2/0 | 14 | 79 | default-switch/bd14 |
| xe-5/0/2 | 15 | 78 | default-switch/bd15 |
| xe-2/2/0 | 15 | 79 | default-switch/bd15 |

### show protection-group ethernet-ring vlan brief

user@host> show protection-group ethernet-ring vlan brief

Ethernet ring IFBD parameters for protection group vkm01

| Interface | Vlan | STP Index | Bridge Domain      |
|-----------|------|-----------|--------------------|
| xe-5/0/2  | 1    | 78        | default-switch/bd1 |
| xe-2/2/0  | 1    | 79        | default-switch/bd1 |
| xe-5/0/2  | 2    | 78        | default-switch/bd2 |
| xe-2/2/0  | 2    | 79        | default-switch/bd2 |
| xe-5/0/2  | 3    | 78        | default-switch/bd3 |
| xe-2/2/0  | 3    | 79        | default-switch/bd3 |
| xe-5/0/2  | 4    | 78        | default-switch/bd4 |
| xe-2/2/0  | 4    | 79        | default-switch/bd4 |
| xe-5/0/2  | 5    | 78        | default-switch/bd5 |
| xe-2/2/0  | 5    | 79        | default-switch/bd5 |
| xe-5/0/2  | 6    | 78        | default-switch/bd6 |
| xe-2/2/0  | 6    | 79        | default-switch/bd6 |
| xe-5/0/2  | 7    | 78        | default-switch/bd7 |
| xe-2/2/0  | 7    | 79        | default-switch/bd7 |
| xe-5/0/2  | 8    | 78        | default-switch/bd8 |
| xe-2/2/0  | 8    | 79        | default-switch/bd8 |
| xe-5/0/2  | 9    | 78        | default-switch/bd9 |

|          |    |    |                     |
|----------|----|----|---------------------|
| xe-2/2/0 | 9  | 79 | default-switch/bd9  |
| xe-5/0/2 | 10 | 78 | default-switch/bd10 |
| xe-2/2/0 | 10 | 79 | default-switch/bd10 |
| xe-5/0/2 | 11 | 78 | default-switch/bd11 |
| xe-2/2/0 | 11 | 79 | default-switch/bd11 |
| xe-5/0/2 | 12 | 78 | default-switch/bd12 |
| xe-2/2/0 | 12 | 79 | default-switch/bd12 |
| xe-5/0/2 | 13 | 78 | default-switch/bd13 |
| xe-2/2/0 | 13 | 79 | default-switch/bd13 |
| xe-5/0/2 | 14 | 78 | default-switch/bd14 |
| xe-2/2/0 | 14 | 79 | default-switch/bd14 |
| xe-5/0/2 | 15 | 78 | default-switch/bd15 |
| xe-2/2/0 | 15 | 79 | default-switch/bd15 |

### show protection-group ethernet-ring vlan detail

user@host> show protection-group ethernet-ring vlan detail

Ethernet ring IFBD parameters for protection group vkm01

Interface name : xe-5/0/2  
 Vlan : 1  
 STP index : 78  
 Bridge Domain : default-switch/bd1

Interface name : xe-2/2/0  
 Vlan : 1  
 STP index : 79  
 Bridge Domain : default-switch/bd1

Interface name : xe-5/0/2  
 Vlan : 2  
 STP index : 78  
 Bridge Domain : default-switch/bd2

Interface name : xe-2/2/0  
 Vlan : 2  
 STP index : 79  
 Bridge Domain : default-switch/bd2

Interface name : xe-5/0/2  
 Vlan : 3  
 STP index : 78  
 Bridge Domain : default-switch/bd3

**show protection-group ethernet-ring vlan group-name vkm01**

```
user@host> show protection-group ethernet-ring vlan vkm01
```

```
Ethernet ring IFBD parameters for protection group vkm01
```

| Interface | Vlan | STP Index | Bridge Domain       |
|-----------|------|-----------|---------------------|
| xe-5/0/2  | 16   | 80        | default-switch/bd16 |
| xe-2/2/0  | 16   | 81        | default-switch/bd16 |
| xe-5/0/2  | 17   | 80        | default-switch/bd17 |
| xe-2/2/0  | 17   | 81        | default-switch/bd17 |
| xe-5/0/2  | 18   | 80        | default-switch/bd18 |
| xe-2/2/0  | 18   | 81        | default-switch/bd18 |
| xe-5/0/2  | 19   | 80        | default-switch/bd19 |
| xe-2/2/0  | 19   | 81        | default-switch/bd19 |
| xe-5/0/2  | 20   | 80        | default-switch/bd20 |
| xe-2/2/0  | 20   | 81        | default-switch/bd20 |
| xe-5/0/2  | 21   | 80        | default-switch/bd21 |
| xe-2/2/0  | 21   | 81        | default-switch/bd21 |
| xe-5/0/2  | 22   | 80        | default-switch/bd22 |
| xe-2/2/0  | 22   | 81        | default-switch/bd22 |
| xe-5/0/2  | 23   | 80        | default-switch/bd23 |
| xe-2/2/0  | 23   | 81        | default-switch/bd23 |
| xe-5/0/2  | 24   | 80        | default-switch/bd24 |
| xe-2/2/0  | 24   | 81        | default-switch/bd24 |
| xe-5/0/2  | 25   | 80        | default-switch/bd25 |
| xe-2/2/0  | 25   | 81        | default-switch/bd25 |
| xe-5/0/2  | 26   | 80        | default-switch/bd26 |
| xe-2/2/0  | 26   | 81        | default-switch/bd26 |
| xe-5/0/2  | 27   | 80        | default-switch/bd27 |
| xe-2/2/0  | 27   | 81        | default-switch/bd27 |
| xe-5/0/2  | 28   | 80        | default-switch/bd28 |
| xe-2/2/0  | 28   | 81        | default-switch/bd28 |
| xe-5/0/2  | 29   | 80        | default-switch/bd29 |
| xe-2/2/0  | 29   | 81        | default-switch/bd29 |
| xe-5/0/2  | 30   | 80        | default-switch/bd30 |
| xe-2/2/0  | 30   | 81        | default-switch/bd30 |

**show protection-group ethernet-ring vlan detail (EX2300 and EX3400 Switches)**

```
user@switch>show protection-group ethernet-ring vlan detail
```

```
Ethernet ring IFBD parameters for protection group pg1001
```

|                |                           |
|----------------|---------------------------|
| Interface name | : ge-0/0/42               |
| Vlan           | : 2001                    |
| STP index      | : 52                      |
| Bridge Domain  | : default-switch/vlan2001 |
| Interface name | : ge-0/0/38               |
| Vlan           | : 2001                    |
| STP index      | : 53                      |
| Bridge Domain  | : default-switch/vlan2001 |

## traceroute ethernet

### Syntax

```
traceroute ethernet
local-mep mep-id
maintenance-association ma-name
maintenance-domain md-name
<ttl value>
<wait seconds>
mac-address | mep-id
<detail>
```

### 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

### 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

**local-mep *mep-id***—(Required when multiple MEPs are configured) Identifier for the local maintenance endpoint.

**detail**—(Optional) Provide detailed information of the responder hostname, ingress port name, egress port name, TTL, and relay action.

**mac-address**—Destination unicast MAC address of the remote maintenance point.

***mep-id***—MEP identifier of the remote maintenance point. The range of values is 1 through 8191.

**maintenance-association *ma-name***—Specifies an existing maintenance association from the set of configured maintenance associations.

**maintenance-domain *md-name***—Specifies an existing maintenance domain from the set of configured maintenance domains.

**ttl *value***—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

**List of Sample Output**  
[traceroute ethernet on page 2046](#)  
[traceroute ethernet detail on page 2047](#)

**Output Fields**  
[Table 196 on page 2045](#) lists the output fields for the **traceroute ethernet** command. Output fields are listed in the approximate order in which they appear.

**Table 196: 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.                                                                                                                                                                                                                                  |
| <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.                                                                                                                                                                                  |

Table 196: traceroute ethernet Output Fields (*continued*)

| Field Name                  | Field Description                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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>                                                                                               |
| <b>Relay Action</b>         | The associated relay action. Relay action can be one of the following: <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         |