

Interfaces User Guide for Switches

Published
2020-06-26

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

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

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

Interfaces User Guide for Switches

Copyright © 2020 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 | xxi

Documentation and Release Notes | xxi

Using the Examples in This Manual | xxi

Merging a Full Example | xxii

Merging a Snippet | xxiii

Documentation Conventions | xxiii

Documentation Feedback | xxvi

Requesting Technical Support | xxvi

Self-Help Online Tools and Resources | xxvii

Creating a Service Request with JTAC | xxvii

1

Configuring Interfaces

Understanding Interfaces | 30

Interfaces Overview for Switches | 30

Network Interfaces for EX Series | 30

Special Interfaces for EX Series | 31

Network Interfaces for EX4600, NFX Series, QFX Series, QFabric System | 34

Special Interfaces for EX4600, NFX Series, QFX Series, QFabric System | 35

Network Interfaces for OCX Series | 36

Special Interfaces for OCX Series | 37

Understanding Interface Naming Conventions | 38

Physical Part of an Interface Name for EX Series | 38

Logical Part of an Interface Name for EX Series | 40

Wildcard Characters in Interface Names for EX Series | 41

Physical Part of an Interface Name for QFX series, NFX Series, EX4600, QFabric System | 41

Logical Part of an Interface Name on a Switch Running QFabric Software Package for QFX series, NFX Series, EX4600, QFabric System | 54

Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software for QFX series, NFX Series, EX4600, QFabric System | 55

Wildcard Characters in Interface Names for QFX series, NFX Series, EX4600, QFabric System | 55

- Physical Part of an Interface Name for OCX1100 | 55
- Wildcard Characters in Interface Names for OCX1100 | 56

Understanding Management Interfaces | 56

Physical Interface Properties | 58

- Damping Shorter Physical Interface Transitions | 59
- Configuring Accounting for the Physical Interface | 60
 - Accounting Profiles Overview | 61
 - Configuring Accounting for the Physical Interface | 61
 - Displaying Accounting Profile for the Physical Interface | 62
- Enabling or Disabling SNMP Notifications on Physical Interfaces | 63
- Configuring Ethernet Loopback Capability | 64
- Configuring Short Reach Mode on QFX5100-48T | 65
- Configuring Flow Control | 66
- Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module | 67
- Setting the Operating Mode on a 2-Port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 Uplink Module | 68
- Configuring the Media Type on Dual-Purpose Uplink Ports | 71
- Disabling a Physical Interface | 72
 - Disabling a Physical Interface | 72
 - Example: Disabling a Physical Interface | 73
 - Effect of Disabling Interfaces on T series PICs | 74

Logical Interface Properties | 75

- Configuring the Interface Address | 76
- Adding a Logical Unit Description to the Configuration | 78
- Configuring the Media MTU | 79
- Setting the Protocol MTU | 80
- Configuring the Interface Bandwidth | 81
- Enabling or Disabling SNMP Notifications on Logical Interfaces | 82
- Configuring Accounting for the Logical Interface | 82
 - Accounting Profiles Overview | 83
 - Configuring Accounting for the Logical Interface | 83
 - Displaying Accounting Profile for the Logical Interface | 84
- Disabling a Logical Interface | 85

Interface Ranges | 87

Understanding Interface Ranges for Switches | 87

Configuring Interface Ranges for EX Series Switches with ELS | 91

Configuring Interface Ranges on Switches | 92

Expanding Interface Range Member and Member Range Statements | 95

Configuration Inheritance for Member Interfaces | 96

Member Interfaces Inheriting Configuration from Configuration Groups | 98

Interfaces Inheriting Common Configuration | 99

Configuring Inheritance Range Priorities | 100

Configuration Expansion Where Interface Range Is Used | 101

Gigabit Ethernet Interface | 102

Speed and Autonegotiation | 102

Configure Interface Speed on Switches | 103

Speed and Autonegotiation | 103

Configure Speed and Autonegotiation on QFX5100-48T Switches | 105

Autonegotiation support for QFX5110-48S and QFX5100-48S with JNP-SFPP-10GE-T transceiver | 106

Autonegotiation support for QFX5120-48Y with JNP-SFPP-10GE-T transceiver | 108

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109

Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 110

Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches | 110

Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches | 111

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches | 112

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches | 112

Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 113

Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support | 114

Configuring VLAN Options and Interface Mode | 114

Configuring the Link Settings | 115

Configuring the IP Options | 118

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches | 119

Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces | 120

Configuring the IP Options | 120

Optical Transport Network (OTN) Interfaces | 121

Understanding the QFX10K-12C-DWDM Line Card | 121

Software Features | 122

OTN Alarms and Defects | 123

Configuring OTN Interface Options on QFX10K-12C-DWDM | 124

Port Settings | 130

Channelize Block of Ports or Individual Port | 131

Interface Naming Conventions | 132

Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches | 134

Channelizing Interfaces on QFX5110-48S Switches | 146

Channelizing Interfaces on QFX5200-32C Switches | 149

Channelizing Interfaces on QFX5210-64C Switches | 152

Channelizing Interfaces on QFX5120-32C Switches | 154

Channelizing Interfaces on QFX5120-48Y Switches | 159

Channelizing Interfaces on QFX5120-48T Switches | 163

Channelizing Interfaces on EX4650-48Y Switches | 165

Understanding Port Ranges and System Modes | 167

Port Ranges for Different Media Types | 168

Supported System Modes | 195

Configuring the System Mode | 197

Configuring the Port Type on QFX3600 Standalone Switches | 202

Configuring the QSFP+ Port Type on QFX3500 Standalone Switches | 204

Configuring the QSFP+ Port Type on QFX5100 Devices | 207

Energy Efficient Ethernet Interfaces | 210

Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces | 210

Configuring Energy Efficient Ethernet on Interfaces | 211

Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port | 211

Disabling EEE on a Base-T Copper Ethernet Port | 211

Verifying That EEE Is Saving Energy on Configured Ports | 212

Uplink Failure Detection | 214

Overview of Uplink Failure Detection | 215

Uplink Failure Detection Configuration | 215

Failure Detection Pair | 216

Debounce Interval | 217

Configuring Interfaces for Uplink Failure Detection | 217

Example: Configuring Interfaces for Uplink Failure Detection | 218

Verifying That Uplink Failure Detection Is Working Correctly | 224

Targeted Broadcast | 225

Understanding Targeted Broadcast | 226

Understanding IP Directed Broadcast | 227

IP Directed Broadcast Overview | 227

IP Directed Broadcast Implementation | 228

When to Enable IP Directed Broadcast | 228

When Not to Enable IP Directed Broadcast | 228

Configuring Targeted Broadcast | 229

Configuring Targeted Broadcast and Its Options | 229

Display Targeted Broadcast Configuration Options | 231

Example: Configuring IP Directed Broadcast on a Switch | 232

Verifying IP Directed Broadcast Status | 237

ARP | 237

Static ARP Table Entries Overview | 238

Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses | 238

Restricted and Unrestricted Proxy ARP Overview | 240

Restricted Proxy ARP | 240

Unrestricted Proxy ARP | 241

Topology Considerations for Unrestricted Proxy ARP | 241

Configuring Restricted and Unrestricted Proxy ARP | 242

Configuring Gratuitous ARP | 243

Resilient Hashing on LAGs and ECMP groups | 245

Understanding the Use of Resilient Hashing to Minimize Flow Remapping in LAGs/ECMP Groups | 245

Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing | 246

Limitations and Caveats for Resilient Hashing | 247

Resilient Hashing on LAGs | 247

Resilient Hashing on ECMP | 248

Configuring Resilient Hashing for LAGs/ECMP Groups | 248

Configuring Resilient Hashing on LAGs | 248

Configuring Resilient Hashing on ECMP Groups | 249

Generic Routing Encapsulation (GRE) | 249

Understanding Generic Routing Encapsulation | 250

Overview of GRE | 250

GRE Tunneling | 251

Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100, QFX10000, and OCX Series Switches | 253

Configuration Limitations | 254

Configuring Generic Routing Encapsulation Tunneling | 255

Configuring a GRE Tunnel | 255

Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly | 256

Understanding Per-Packet Load Balancing | 257

Configuring Aggregated Ethernet Interfaces

Aggregated Ethernet Interfaces | 261

Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262

Link Aggregation Group | 263

Link Aggregation Control Protocol (LACP) | 265

Forcing LAG Links or Interfaces with Limited LACP Capability to Be Up | 267

Configuring an Aggregated Ethernet Interface | 267

Configuring Tagged Aggregated Ethernet Interfaces | 268

Configuring Untagged Aggregated Ethernet Interfaces | 270

Configuring the Number of Aggregated Ethernet Interfaces on the Device (Enhanced Layer 2 Software) | 271

Example: Configuring Aggregated Ethernet Interfaces | 271

Deleting an Aggregated Ethernet Interface | 273

Understanding Local Link Bias | 274

Configuring Local Link Bias | 276

Understanding Local Minimum Links | 276

Configuring Local Minimum Links | 278

Local Minimum Links Effect on LAG Minimum Links | 279

Local Minimum Links and Local Link Bias | 279

Troubleshooting an Aggregated Ethernet Interface | 279

Show Interfaces Command Shows the LAG is Down | 280

Logical Interface Statistics Do Not Reflect All Traffic | 280

IPv6 Interface Traffic Statistics Are Not Supported | 280

SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0 | 280

Configuring Link Aggregation | 281

Creating an Aggregated Ethernet Interface | 282

Configuring the VLAN Name and VLAN ID Number | 283

Configuring Aggregated Ethernet LACP (CLI Procedure) | 283

Configuring Aggregated Ethernet Link Protection | 286

Configuring Link Protection for Aggregated Ethernet Interfaces | 286

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces | 287

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link | 287

Disabling Link Protection for Aggregated Ethernet Interfaces | 288

Configuring Aggregated Ethernet Link Speed | 288

Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface | 290

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291

Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch | 299

Configuring Aggregated Ethernet LACP | 305

Configuring the LACP Interval | 307

Configuring LACP Link Protection | 307

Configuring LACP System Priority | 309

Configuring LACP System Identifier | 309

Configuring LACP administrative Key	310
Configuring LACP Port Priority	310
Tracing LACP Operations	311
LACP Limitations	311
Example: Configuring Aggregated Ethernet LACP	311
Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches	313
Configuring LACP Link Protection for a Single Link at the Global Level	315
Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level	315
Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface	316
Configuring LACP Hold-UP Timer to Prevent Link Flapping on LAG Interfaces	319
Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets	320
Verifying the LACP Setup	320
Verifying That LACP Packets Are Being Exchanged	320
Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch	321
Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch	328
Understanding Independent Micro BFD Sessions for LAG	333
Configuring Micro BFD Sessions for LAG	336
Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic	342
Understanding the Hashing Algorithm	343
IP (IPv4 and IPv6)	344
MPLS	346
MAC-in-MAC Packet Hashing	347
Layer 2 Header Hashing	348
Hashing Parameters	349
Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)	349
Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing	350
Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing	350
Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing	351

Configuring Other Hashing Parameters | 351

Load Balancing for Aggregated Ethernet Interfaces | 352

Load Balancing and Ethernet Link Aggregation Overview | 353

Configuring Load Balancing Based on MAC Addresses | 353

Configuring Load Balancing on a LAG Link | 355

Example: Configuring Load Balancing on a LAG Link | 355

Understanding Consistent Load Balancing Through Resilient Hashing on ECMP Groups | 356

Configuring Consistent Load Balancing for ECMP Groups | 357

Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches | 360

Create LAGs for Multicasting in Increments of 10 Gigabits | 361

When Should I Use Multicast Load Balancing? | 362

How Does Multicast Load Balancing Work? | 363

How Do I Implement Multicast Load Balancing on an EX8200 Switch? | 364

Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches | 365

Dynamic Load Balancing | 372

Configuring Dynamic Load Balancing | 374

Example: Configure Dynamic Load Balancing | 376

3

Flexible Ethernet Services Encapsulation

Flexible Ethernet Services Encapsulation | 385

Understanding Flexible Ethernet Services Encapsulation on Switches | 385

Service Provider Style | 386

Enterprise Style | 386

Flexible Ethernet Services | 387

Configuring Flexible Ethernet Services Encapsulation to Support the Service Provider and Enterprise Styles of Configuration | 388

4

Monitoring and Troubleshooting Information

Monitoring Interfaces | 392

Monitoring Interface Status and Traffic | 392

Monitoring System Process Information | 392

Monitoring System Properties | 393

Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 395

Tracing Operations of the Interface Process | 397

Troubleshooting Interfaces | 399

Troubleshooting Network Interfaces | 399

Statistics for logical interfaces on Layer 2 interfaces are not accurate | 399

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down | 400

Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) | 400

Troubleshooting Uplink Ports on EX2300 Switches | 403

Speeds 10-Mbps and 100-Mbps not supported on uplink ports 4 and 5 on EX2300-48MP switches | 404

1

Configuration Statements and Operational Commands

Configuration Statements: Interfaces | 407

address | 409

auto-negotiation | 413

autostate-exclude | 414

bandwidth (Interfaces) | 416

broadcast | 418

ccc | 419

configured-flow-control | 420

description (Interfaces) | 422

disable (Interface) | 424

ethernet (Alarm) | 426

ethernet-switching | 427

eui-64 | 428

family | 429

forward-and-send-to-re | 434

forward-only | 435

filter | 436

hold-time (Physical Interface) | 438

inet (interfaces) | 441

inet6 (interfaces) | 442

inet (enhanced-hash-key) | 443

- inet6 (enhanced-hash-key) | 446
- interface (Multichassis Protection) | 448
- interface-mode | 449
- interface-range | 451
- interfaces (QFX Series) | 453
- interfaces (EX Series switches) | 464
- irb (Interfaces) | 475
- loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | 480
- mac | 482
- media-type (Dual-Purpose Uplink Ports) | 483
- member | 484
- member-range | 485
- mode (Interfaces) | 486
- mtu | 487
- nd6-stale-time | 489
- no-redirects | 490
- policer (MAC) | 491
- preferred | 493
- primary (Address on Interface) | 494
- traceoptions (Individual Interfaces) | 495
- reflective-relay | 504
- speed (Ethernet) | 505
- traps | 513
- unidirectional | 515
- unit | 516

Configuration Statements: Gigabit Ethernet Interfaces | 518

- container-devices | 519
- craft-lockout | 520
- no-neighbor-learn | 522

Configuration Statements: OTN Interfaces | 524

- alarm low-light-alarm | 525
- encoding | 526

fec | 527
laser-enable | 528
line-loopback | 529
link-event-rate | 530
modulation-format | 531
optics-options | 532
prbs | 533
preemptive-fast-reroute | 534
signal-degrade | 535
tca | 536
trigger | 538
tx-power | 540
warning | 541
wavelength | 542

Configuration Statements: Aggregated Ethernet Interfaces | 547

802.3ad | 549
alarm (chassis) | 551
aggregated-devices | 553
aggregated-ether-options | 555
backup-liveness-detection | 559
backup-peer-ip | 560
bfd-liveness-detection (LAG) | 561
chassis (EX Series) | 564
chassis (QFabric System) | 566
chassis (QFX Series) | 569
device-count | 571
disable (Link Protection) | 572
disk-failure-action | 573
disable (Multicast Load Balancing) | 574
dlb | 575
ether-options | 577
ether-type | 585
ethernet | 586

ethernet (Aggregated Devices) | 587

fibre-channel (Alarm) | 588

flow-control | 589

force-up | 591

hash-parameters | 592

interconnect-device (Chassis) | 594

iccp | 596

lACP (802.3ad) | 598

lACP (Aggregated Ethernet) | 600

link-down | 603

link-mode | 604

link-protection | 606

link-protection-sub-group (802.3ad) | 608

link-protection-sub-group (aggregated-ether-options) | 609

link-speed | 611

liveness-detection | 614

local-bias | 615

local-ip-addr (ICCP) | 616

local-minimum-links-threshold | 617

management-ethernet (Alarm) | 619

minimum-interval (Liveness Detection) | 620

minimum-links | 621

minimum-receive-interval (Liveness Detection) | 623

multicast-loadbalance | 624

multiservice | 626

node-device (Chassis) | 628

node-group (Chassis) | 629

non-revertive (Chassis) | 631

non-revertive (Interfaces) | 632

on-disk-failure | 633

optics-options | 634

peer (ICCP) | 636

periodic | 638

port-priority | 640

routing-engine | 641

rx-buffers | 642

session-establishment-hold-time | 644

transmit-interval (Liveness Detection) | 645

tx-buffers | 646

Configuration Statements: Channelizing Interfaces | 648

channel-speed | 649

fpc | 650

fte (Port) | 652

number-of-sub-ports | 653

pic | 655

pic-mode | 658

sfpplus | 659

short-reach-mode | 660

xe (Port) | 661

xle (Port) | 663

Configuration Statements: Energy Efficient Interfaces | 665

ieee-802-3az-eee | 665

Configuration Statements: VLANs | 666

ethernet-switch-profile | 667

l2-domain-id-for-l3 | 669

layer2 (enhanced-hash-key) | 670

layer3-domain-identifier | 672

members | 673

no-local-switching | 675

port-mode | 676

tag-protocol-id (TPIDs Expected to Be Sent or Received) | 678

vlan-id | 680

vlan-tagging | 681

Configuration Statements: Uplink Failure Detection for Interfaces | 684

debounce-interval | 685
group | 686
link-to-disable | 687
link-to-monitor | 688
uplink-failure-detection | 689

Configuration Statements: Unicast Reverse Path Forwarding (uRPF) | 690

group (RPF Selection) | 691
multicast-rpf-routes | 692
next-hop (PIM RPF Selection) | 693
prefix-list (PIM RPF Selection) | 694
rpf-check (Dynamic Profiles) | 695
rpf-check | 696
rpf-check-policy (Routing Options RPF) | 698
rpf-loose-mode-discard | 699
rpf-selection | 700
source (PIM RPF Selection) | 702
unicast-reverse-path | 703
wildcard-source (PIM RPF Selection) | 704

Configuration Statements: IP Directed and Targeted Broadcast | 705

targeted-broadcast | 706
policy-statement | 708

Configuration Statements: ARP | 714

arp (Interfaces) | 715
gratuitous-arp-reply | 719
no-gratuitous-arp-request | 720
proxy-arp | 721

Configuration Statements: Resilient Hashing | 723

ecmp-resilient-hash | 724
enhanced-hash-key | 725
hash-key (Forwarding Options) | 731

hash-mode | 733

hash-seed | 735

inet (enhanced-hash-key) | 736

inet6 (enhanced-hash-key) | 739

ipv6-flow-label | 742

resilient-hash | 743

Configuration Statements: Generic Routing Encapsulation (GRE) | 744

gtp-header-offset | 745

gtp-tunnel-endpoint-identifier | 747

source | 749

ttl | 750

tunnel | 751

allow-fragmentation | 752

copy-tos-to-outer-ip-header | 753

do-not-fragment | 754

destination (Tunnels) | 755

family | 756

routing-instance | 758

source | 759

tunnel | 760

tunnel-port | 761

unit (Interfaces) | 762

Configuration Statements: Flexible Ethernet Services Encapsulation | 764

encapsulation | 765

encapsulation (Logical Interface) | 772

flexible-vlan-tagging | 777

Operational Commands | 778

Common Output Fields Description | 779

Damping Field | 779

Destination Class Field | 779

Enabled Field | 780

Filters Field | 780

Flags Fields | 781**Addresses, Flags Field | 781****Device Flags Field | 782****Family Flags Field | 782****Interface Flags Field | 783****Link Flags Field | 784****Logical Interface Flags Field | 784****Label-Switched Interface Traffic Statistics Field | 785****Policer Field | 786****Protocol Field | 786****RPF Failures Field | 787****Source Class Field | 787****clear interfaces statistics | 788****monitor interface | 790****request diagnostics tdr | 806****request chassis system-mode | 808****Show chassis system-mode | 810****show diagnostics tdr | 812****show forwarding-options enhanced-hash-key | 818****show interfaces (Discard) | 825****show interfaces | 832****show interfaces (Serial) | 941****show interfaces diagnostics optics | 959****show interfaces extensive | 990****show interfaces fabric | 1047****show interfaces ge | 1075****show interfaces (GRE) | 1092****show interfaces irb | 1104****show interfaces mc-ae | 1113****show interfaces me0 | 1117****show interfaces queue | 1127****show interfaces queue fabric | 1183****show interfaces xe | 1213****show interfaces xle | 1237**

[show interfaces statistics fabric | 1260](#)
[show interfaces vlan | 1285](#)
[show lacp interfaces | 1301](#)
[show lacp statistics interfaces \(View\) | 1308](#)
[show redundant-trunk-group | 1310](#)
[show uplink-failure-detection | 1312](#)
[show virtual-chassis vc-port diagnostics optics | 1315](#)
[test interface restart-auto-negotiation | 1334](#)

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | xxi
- Using the Examples in This Manual | xxi
- Documentation Conventions | xxiii
- Documentation Feedback | xxvi
- Requesting Technical Support | xxvi

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

Documentation and Release Notes

To obtain the most current version of all Juniper Networks[®] technical documentation, see the product documentation page on the Juniper Networks website at <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 xxiv](#) 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 xxiv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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


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

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

Documentation Feedback

We encourage you to provide feedback 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. 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

CHAPTER

Configuring Interfaces

Understanding Interfaces | 30

Physical Interface Properties | 58

Logical Interface Properties | 75

Interface Ranges | 87

Gigabit Ethernet Interface | 102

Optical Transport Network (OTN) Interfaces | 121

Port Settings | 130

Energy Efficient Ethernet Interfaces | 210

Uplink Failure Detection | 214

Targeted Broadcast | 225

ARP | 237

Resilient Hashing on LAGs and ECMP groups | 245

Generic Routing Encapsulation (GRE) | 249

Understanding Per-Packet Load Balancing | 257

Understanding Interfaces

IN THIS SECTION

- [Interfaces Overview for Switches | 30](#)
- [Understanding Interface Naming Conventions | 38](#)
- [Understanding Management Interfaces | 56](#)

Junos OS supports different types of interfaces on which the devices function. The following topics provide information of types of interfaces used, the naming conventions and the usage of management interfaces by Juniper Networks.

Interfaces Overview for Switches

IN THIS SECTION

- [Network Interfaces for EX Series | 30](#)
- [Special Interfaces for EX Series | 31](#)
- [Network Interfaces for EX4600, NFX Series, QFX Series, QFabric System | 34](#)
- [Special Interfaces for EX4600, NFX Series, QFX Series, QFabric System | 35](#)
- [Network Interfaces for OCX Series | 36](#)
- [Special Interfaces for OCX Series | 37](#)

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

Network Interfaces for EX Series

Network interfaces connect to the network and carry network traffic. [Table 3 on page 31](#) lists the types of network interfaces supported on EX Series switches.

Table 3: Network Interfaces Types and Purposes for EX Series

Type	Purpose
Aggregated Ethernet interfaces	<p>All EX Series switches allow you to group Ethernet interfaces at the physical layer to form a single link layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i>. These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.</p> <p>See “Understanding Aggregated Ethernet Interfaces and LACP for Switches” on page 262.</p>
LAN access interfaces	<p>Use these EX Series switch interfaces to connect a personal computer, laptop, file server, or printer to the network. When you power on an EX Series switch and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports. The default configuration also enables autonegotiation for both speed and link mode.</p>
Power over Ethernet (PoE) interfaces	<p>EX Series switches provide PoE network ports with various switch models. These ports can be used to connect voice over IP (VoIP) telephones, wireless access points, video cameras, and point-of-sale devices to safely receive power from the same access ports that are used to connect personal computers to the network. PoE interfaces are enabled by default in the factory configuration.</p> <p>See <i>Understanding PoE on EX Series Switches</i>.</p>
Trunk interfaces	<p>EX Series access switches can be connected to a distribution switch or customer-edge (CE) switches or routers. To use a port for this type of connection, you must explicitly configure the network interface for trunk mode. The interfaces from the distribution switch or CE switch to the access switches must also be configured for trunk mode.</p>

Special Interfaces for EX Series

[Table 4 on page 31](#) lists the types of special interfaces supported on EX Series switches.

Table 4: Special Interfaces Types and Purposes for EX Series

Type	Purpose
Console port	<p>Each EX Series switch has a serial port, labeled CON or CONSOLE, for connecting tty-type terminals to the switch using standard PC-type tty cables. The console port does not have a physical address or IP address associated with it. However, it is an interface since it provides access to the switch. On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, or an EX4500 Virtual Chassis, you can access the master and configure all members of the Virtual Chassis through any member's console port. For more information about the console port in a Virtual Chassis, see <i>Understanding Global Management of a Virtual Chassis</i>.</p>

Table 4: Special Interfaces Types and Purposes for EX Series (*continued*)

Type	Purpose
Loopback	All EX Series switches have this software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	<p>The Juniper Networks Junos operating system (Junos OS) for EX Series switches automatically creates the switch's management Ethernet interface, me0. The management Ethernet interface provides an out-of-band method for connecting to the switch. To use me0 as a management port, you must configure its logical port, me0.0, with a valid IP address. You can connect to the management interface over the network using utilities such as SSH or Telnet. SNMP can use the management interface to gather statistics from the switch. (The management interface me0 is analogous to the fxp0 interfaces on routers running Junos OS.)</p> <p>See “Understanding Management Interfaces” on page 56.</p>
Integrated Routing and Bridging (IRB) Interface or Routed VLAN Interface (RVI)	<p>EX Series switches use an integrated routing and bridging (IRB) interface or Routed VLAN Interface (RVI) to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The IRB interface or RVI functions as a logical router, eliminating the need for having both a switch and a router. These interfaces must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed from.</p> <p>See <i>Understanding Integrated Routing and Bridging</i>.</p>

Table 4: Special Interfaces Types and Purposes for EX Series (*continued*)

Type	Purpose
Virtual Chassis port (VCP) interfaces	<p>Virtual Chassis ports (VCPs) are used to interconnect switches in a Virtual Chassis:</p> <ul style="list-style-type: none"> EX3300 switches—Port 2 and port 3 of the SFP+ uplink ports are preconfigured as VCPs and can be used to interconnect up to six EX3300 switches in an EX3300 Virtual Chassis. See <i>Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port</i>. EX4200 and EX4500 switches—Each EX4200 switch or each EX4500 switch with a Virtual Chassis module installed has two dedicated VCPs on its rear panel. These ports can be used to interconnect up to ten EX4200 switches in an EX4200 Virtual Chassis, up to ten EX4500 switches in an EX4500 Virtual Chassis, and up to ten switches in a mixed EX4200 and EX4500 Virtual Chassis. When you power on switches that are interconnected in this manner, the software automatically configures the VCP interfaces for the dedicated ports that have been interconnected. These VCP interfaces are not configurable or modifiable. See <i>Understanding the High-Speed Interconnection of the Dedicated Virtual Chassis Ports Connecting EX4200, EX4500, and EX4550 Member Switches</i>. <p>You can also interconnect EX4200 and EX4500 switches by using uplink module ports. Using uplink ports allows you to connect switches over longer distances than you can by using the dedicated VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See <i>Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port</i>.</p> <ul style="list-style-type: none"> EX4300 switches—All QSFP+ ports are configured as VCPs by default. See <i>Understanding EX Series Virtual Chassis</i>. <p>You can also interconnect EX4300 switches into a Virtual Chassis by using SFP+ uplink module ports as VCPs. Using uplink ports as VCPs allows you to connect switches over longer distances than you can by using the QSFP+ ports as VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See <i>Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port</i>.</p> <ul style="list-style-type: none"> EX8200 switches—EX8200 switches can be connected to an XRE200 External Routing Engine to create an EX8200 Virtual Chassis. The XRE200 External Routing Engine has dedicated VCPs that connect to ports on the internal Routing Engines of the EX8200 switches and can connect to another XRE200 External Routing Engine for redundancy. These ports require no configuration. <p>You can also connect two members of an EX8200 Virtual Chassis so that they can exchange Virtual Chassis Control Protocol (VCCP) traffic. To do so, you explicitly configure network ports on the EX8200 switches as VCPs.</p>

Table 4: Special Interfaces Types and Purposes for EX Series (*continued*)

Type	Purpose
Virtual management Ethernet (VME) interface	<p>EX3300, EX4200, EX4300, and EX4500 switches have a VME interface. This is a logical interface that is used for Virtual Chassis configurations and allows you to manage all the members of the Virtual Chassis through the master. For more information about the VME interface, see <i>Understanding Global Management of a Virtual Chassis</i>.</p> <p>EX8200 switches do not use a VME interface. An EX8200 Virtual Chassis is managed through the management Ethernet (me0) interface on the XRE200 External Routing Engine.</p>

Network Interfaces for EX4600, NFX Series, QFX Series, QFabric System

Network interfaces connect to the network and carry network traffic. [Table 5 on page 34](#) lists the types of network interfaces supported.

Table 5: Network Interfaces Types and Purposes for EX4600, NFX Series, QFX Series, QFabric System

Type	Purpose
Aggregated Ethernet interfaces	Group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.
Channelized Interfaces	<p>Depending on the device and software package, 40-Gbps QSFP+ ports can be configured to operate as the following types of interfaces:</p> <ul style="list-style-type: none"> • 10-Gigabit Ethernet interfaces (<i>xe</i>) • 40-Gigabit Ethernet interfaces (<i>et</i> and <i>xle</i>) • 40-Gigabit data plane uplink interfaces (<i>fte</i>) <p>When an <i>et</i> port is channelized to four <i>xe</i> ports, a colon is used to signify the four separate channels. For example, on a QFX3500 standalone switch with port 2 on PIC 1 configured as four 10-Gigabit Ethernet ports, the interface names are <i>xe-0/1/2:0</i>, <i>xe-0/1/2:1</i>, <i>xe-0/1/2:2</i>, and <i>xe-0/1/2:3</i></p> <p>NOTE: You cannot configure channelized interfaces to operate as Virtual Chassis ports.</p>
Ethernet Interfaces	Configure Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet interfaces to connect to other servers, storage, and switches. You can configure 40-Gigabit data plane uplink ports to connect a Node device to an Interconnect devices as well as for Virtual Chassis ports (VCPs).

Table 5: Network Interfaces Types and Purposes for EX4600, NFX Series, QFX Series, QFabric System (continued)

Type	Purpose
Fibre Channel interfaces	<p>Use Fibre Channel interfaces to connect the switch to a Fibre Channel over Ethernet (FCoE) forwarder or a Fibre Channel switch in a storage area network (SAN). You can configure Fibre Channel interfaces only on ports 0 through 5 and 42 through 47 on QFX3500 devices. Fibre Channel interfaces do not forward Ethernet traffic.</p> <p><i>See Overview of Fibre Channel.</i></p>
LAN access interfaces	<p>Use these interfaces to connect to other servers, storage, and switches. When you power on a QFX Series product and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports.</p>
Multichassis aggregated Ethernet (MC-AE) interfaces	<p>Group a LAG on one standalone switch with a LAG on another standalone switch to create a MC-AE. The MC-AE provides load balancing and redundancy across the two standalone switches.</p>
Tagged-access mode interfaces	<p>Use tagged-access interfaces to connect a switch to an access layer device. Tagged-access interfaces can accept VLAN-tagged packets from multiple VLANs.</p>
Trunk interfaces	<p>Use trunk interfaces to connect to other switches or routers. To use a port for this type of connection, you must explicitly configure the port interface for trunk mode. The interfaces from the switches or routers must also be configured for trunk mode. In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN.</p>
Virtual Chassis ports (VCPs)	<p>You can use Virtual Chassis ports to send and receive Virtual Chassis Control Protocol (VCCP) traffic, and to create, monitor, and maintain the Virtual Chassis. On QFX3500, QFX3600, QFX5100, QFX5110, QFX5200, and EX4600 standalone switches, you can configure 40-Gigabit Ethernet QSFP+ uplink ports (non-channelized) or fixed SFP+ 10-Gigabit Ethernet ports as VCPs by issuing the request virtual-chassis-vc-port-set CLI command. QFX5110 switches also support configuring 100-Gigabit QSFP28 ports as VCPs.</p>

Special Interfaces for EX4600, NFX Series, QFX Series, QFabric System

Table 6 on page 36 lists the types of special interfaces supported.

Table 6: Special Interfaces Types and Purposes supported on EX4600, NFX Series, QFX Series, QFabric System

Type	Purpose
Console port	Each device has a serial console port, labeled CON or CONSOLE , for connecting tty-type terminals to the switch. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch.
Loopback interface	A software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	<p>The management Ethernet interface provides an out-of-band method for connecting to a standalone switch and QFabric system.</p> <p>NOTE: On OCX Series switches, the em0 management interface always has the status up in show command outputs, even if the physical port is empty. The me0 interface is a virtual interface between Junos and the host operating system, therefore its status is independent from the status of the physical port.</p>
Routed VLAN interfaces (RVI and IRB interfaces)	<p>Layer 3 routed VLAN interfaces (called RVI in the original CLI, and called IRB in Enhanced Layer 2 Software) route traffic from one broadcast domain to another and perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The RVI or IRB functions as a logical router, eliminating the need for having both a switch and a router. The RVI or IRB must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it.</p>

Network Interfaces for OCX Series

Network interfaces connect to the network and carry network traffic. [Table 7 on page 36](#) lists the types of network interfaces supported.

Table 7: Network Interfaces Types and Purposes for OCX Series

Type	Purpose
Aggregated Ethernet interfaces	Group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.
Ethernet Interfaces	Configure Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet interfaces to connect to other servers, storage, and switches.

Special Interfaces for OCX Series

Table 8 on page 37 lists the types of special interfaces supported.

Table 8: Special Interfaces Types and Purposes for OCX Series

Type	Purpose
Console port	Each device has a serial console port, labeled CON or CONSOLE , for connecting tty-type terminals to the switch. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch.
Loopback interface	A software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	<p>The management Ethernet interface provides an out-of-band method for connecting to a standalone switch and QFabric system.</p> <p>NOTE: On OCX Series switches, the em0 management interface always has the status up in show command outputs, even if the physical port is empty. The me0 interface is a virtual interface between Junos and the host operating system, therefore its status is independent from the status of the physical port.</p>

SEE ALSO

<i>EX2200 Switches Hardware Overview</i>
<i>EX3200 System Overview</i>
<i>EX3300 Switches Hardware Overview</i>
<i>EX4200 Switches Hardware Overview</i>
<i>EX4300 Switches Hardware Overview</i>
<i>EX4500 Switches Hardware Overview</i>
<i>EX6210 Switch Hardware Overview</i>
<i>EX8208 Switch Hardware Overview</i>
<i>EX8216 Switch Hardware Overview</i>
<i>Understanding Layer 3 Logical Interfaces</i>
<i>Understanding Layer 3 Subinterfaces</i>

Understanding Interface Naming Conventions

IN THIS SECTION

- Physical Part of an Interface Name for EX Series | 38
- Logical Part of an Interface Name for EX Series | 40
- Wildcard Characters in Interface Names for EX Series | 41
- Physical Part of an Interface Name for QFX series, NFX Series, EX4600, QFabric System | 41
- Logical Part of an Interface Name on a Switch Running QFabric Software Package for QFX series, NFX Series, EX4600, QFabric System | 54
- Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software for QFX series, NFX Series, EX4600, QFabric System | 55
- Wildcard Characters in Interface Names for QFX series, NFX Series, EX4600, QFabric System | 55
- Physical Part of an Interface Name for OCX1100 | 55
- Wildcard Characters in Interface Names for OCX1100 | 56

The EX Series, QFX Series, NFX Series, OCX1100, QFabric System, and EX4600 devices use a naming convention for defining the interfaces that are similar to that of other platforms running under Juniper Networks Junos OS. This topic provides brief information about the naming conventions used for interfaces on the QFX Series and on EX4600 switches.

For detailed information on interface naming like physical part, logical part, and channel part of the interfaces, see *Interface Naming Overview*.

This topic describes:

Physical Part of an Interface Name for EX Series

Network interfaces in Junos OS are specified as follows:

type-fpc / pic / port

EX Series switches apply this convention as follows:

- *type*—EX Series interfaces use the following media types:
 - *ge*—Gigabit Ethernet interface
 - *xe*—10 Gigabit Ethernet interface

- et—40 Gigabit Ethernet interface
- *fpc*—Flexible PIC Concentrator. EX Series interfaces use the following convention for the FPC number in interface names:
 - On an EX2200 switch, an EX2300, an EX3200 switch, a standalone EX3300 switch, a standalone EX3400 switch, a standalone EX4200 switch, a standalone EX4300 switch, a standalone EX4500, and a standalone EX4550 switch, FPC refers to the switch itself. The FPC number is **0** by default on these switches.
 - On an EX3300 Virtual Chassis, an EX3400 Virtual Chassis, an EX4200 Virtual Chassis, an EX4300 Virtual Chassis, an EX4500 Virtual Chassis, an EX4550 Virtual Chassis, or a mixed Virtual Chassis, the FPC number indicates the member ID of the switch in the Virtual Chassis.
 - On an EX6200 switch and a standalone EX8200 switch, the FPC number indicates the slot number of the line card that contains the physical interface. On an EX6200 switch, the FPC number also indicates the slot number of the Switch Fabric and Routing Engine (SRE) module that contains the uplink port.
 - On an EX8200 Virtual Chassis, the FPC number indicates the slot number of the line card on the Virtual Chassis. The line card slots on Virtual Chassis member 0 are numbered 0 through 15; on member 1, they are numbered 16 through 31, and so on.
 - On EX9251 switch, the FPC number is always **0**.
 - The EX9253 switch does not have actual FPCs—the line cards are the FPC equivalents on the switch. In FPC (n), n is a value in the range of 0-1. The value corresponds to the line card slot number in which the line card is installed.
 - On an EX29204 switch, switch does not have actual FPCs—the line cards are the FPC equivalents on the switch. The value ranges from 0-2, and it corresponds to the line card slot number in which the line card is installed.
- *pic*—EX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
 - On EX2200, EX2300, EX3200, EX3300, EX4200, EX4500 switch, and EX4550 switches, the PIC number is **0** for all built-in interfaces (interfaces that are not uplink ports).
 - On EX2200, EX2300, EX3200, EX3300, and EX4200 switches, the PIC number is **1** for uplink ports.
 - On EX3400 switches, the PIC number is **0** for built-in network ports, **1** for built-in QSFP+ ports (located on the rear panel of the switch), and **2** for uplink module ports.
 - On EX4300 switches, the PIC number is **0** for built-in network ports, **1** for built-in QSFP+ ports (located on the rear panel of the switch), and **2** for uplink module ports.
 - On EX4500 switches, the PIC number is **1** for ports on the left-hand uplink module and **2** for ports on the right-hand uplink module.

- On EX4550 switches, the PIC number is **1** for ports in the expansion module or Virtual Chassis module installed in the module slot on the front panel of the switch and **2** for those in the expansion module or Virtual Chassis module installed in the module slot on the rear panel of the switch.
- On EX6200 and EX8200 switches, the PIC number is always **0**.
- On EX9251 and EX9253 switches, the PIC number is **0** for built-in network ports, **1** for built-in QSFP+ ports (located on the rear panel of the switch).
- On EX9204 switches, the PIC number ranges from 0-3.
- *port*—EX Series interfaces use the following convention for port numbers:
 - On EX2200, EX2300, EX3200, EX3300, EX3400, EX4200, EX4300, EX4500, and EX4550 switches, built-in network ports are numbered from left to right. On models that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.
 - Uplink ports in EX2200, EX3200, EX3300, EX3400, EX4200, EX4300, EX4500, and EX4550 switches are labeled from left to right, starting with **0**.
 - On EX6200 and EX8200 switches, the network ports are numbered from left to right on each line card. On line cards that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.
 - Uplink ports on an SRE module in an EX6200 switch are labeled from left to right, starting with **0**.
 - EX9251 Switch has eight 10-Gigabit Ethernet ports and four rate-selectable ports that you can configure as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet ports; each rate-selectable port can be configured as four 10-Gigabit Ethernet ports by using a breakout cable. The 10-Gigabit Ethernet ports support SFP+ transceivers and rate-selectable ports support QSFP28 and QSFP+ transceivers.
 - EX9253 contains six built-in QSFP+ ports, each of which can house QSFP+ pluggable transceivers and 12 built-in QSFP28 ports, each of which can house QSFP28 pluggable transceivers.

Logical Part of an Interface Name for EX Series

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

Interface	State	VLAN members	Blocking
ge-0/0/0.0	down	remote-analyzer	unblocked

ge-0/0/1.0	down	default	unblocked
ge-0/0/10.0	down	default	unblocked

Wildcard Characters in Interface Names for EX Series

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

Physical Part of an Interface Name for QFX series, NFX Series, EX4600, QFabric System

Interfaces in Junos OS are specified as follows:

device-name:type-fpc/pic/port

The convention is as follows (and platform support depends on the Junos OS release in your installation):

- *device-name*—(QFabric systems only) The *device-name* is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.
- *type*—The QFX Series and EX4600 device interfaces use the following media types:
 - **fc**—Fibre Channel interface
 - **ge**—Gigabit Ethernet interface
 - **xe**—10-Gigabit Ethernet interface
 - **sxe**—10-Gigabit Service interface. **sxe** is an internal interface and user must not configure this interface. It supports L2 and L3 configurations like VLANs and IP address.
 - **xle**—40-Gigabit Ethernet interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
 - **et**—25-Gigabit Ethernet interface (QFX5120 and QFX5200 switches)
 - **et**—40-Gigabit Ethernet interface (QFX3500, QFX3600, QFX5100, QFX5200, QFX10000, and EX4600 switches running Enhanced Layer 2 Software)
 - **et**—100-Gigabit Ethernet interface (QFX5200 and QFX10000 switches running Enhanced Layer 2 Software)
 - **fte**—40-Gigabit data plane uplink interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
 - **me**—Management interface
 - **em**—Management interface on QFX5100 and EX4600 switches.

- *fpc*—Flexible PIC Concentrator. QFX Series interfaces use the following convention for the FPC number in interface names:

- On QFX3500, QFX3600, QFX5100 devices running a QFabric software package, and QFX10002 switches, the FPC number is always 0.

The FPC number indicates the slot number of the line card that contains the physical interface.

- On QFX3500, QFX3600, QFX5100, QFX5200, EX4600, QFX10002, QFX10008, and QFX10016 switches running Enhanced Layer 2 Software, the member ID of a member in a Virtual Chassis determines the FPC number.

NOTE: Every member in a Virtual Chassis must have a unique member ID, otherwise the Virtual Chassis will not be created.

- On standalone QFX5100, EX4600, and QFX10002 switches, the FPC number is always 0.
- *pic*—QFX Series and EX4600 device interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:

Table 9: Naming Conventions for PICs

Device with Software Package	Convention
QFX3500 switch with QFabric software package	PIC 0 can support 48 ports, PIC 1 can support 16 10-Gigabit Ethernet ports, and PIC 2 can support 4 40-Gigabit Ethernet ports.
QFX3500 switch with Enhanced Layer 2 software	PIC 0 can support 48 ports, and PIC 1 can support 16 10-Gigabit Ethernet ports, and 4 40-Gigabit Ethernet ports.
QFX3500 Node device with a QFabric software package	PIC 0 can support 48 ports and PIC 1 can support four 40-Gigabit data plane uplink ports.
QFX3600 switch with a QFabric software package	PIC 0 can support 64 10-Gigabit Ethernet ports, and PIC 1 can support 16 40-Gigabit Ethernet ports.
QFX3600 switch with Enhanced Layer 2 software	PIC 0 can support 64 10-Gigabit Ethernet ports and can also support 16 40-Gigabit Ethernet ports.
QFX3600 Node device running a QFabric software package	PIC 0 can support 56 10-Gigabit Ethernet ports, and PIC 1 can support 8 40-Gigabit data plane uplink ports, and up to 14 40-Gigabit Ethernet ports.
QFX5100-48S switch with Enhanced Layer 2 software	PIC 0 provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces.

Table 9: Naming Conventions for PICs (*continued*)

Device with Software Package	Convention
EX4600 device with Enhanced Layer 2 software	PIC 0 provides 4 40-Gbps QSFP+ ports and 24 10-Gigabit Ethernet interfaces. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provide 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 10-Gbps SFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each.
QFX5100-48S switch with a QFabric software package	PIC 1 provides six 40-Gbps QSFP+ ports, and PIC 0 provides 48 10-Gigabit Ethernet interfaces.
QFX5100-24Q switch with Enhanced Layer 2 software	PIC 0 provides 24 40-Gbps QSFP+ ports. PIC 1 and PIC 2 can each contain a QFX-EM-4Q expansion module, and each expansion module provides 4 40-Gbps QSFP+ ports
QFX5100-96S switch with Enhanced Layer 2 software	PIC 0 provides 96 10-Gigabit Ethernet interfaces and 8 40-Gbps QSFP+ ports .
QFX5110-48S switch with Enhanced Layer 2 software	PIC 0 can support 48 10-Gigabit Ethernet ports labeled 0 through 47, and 4 QSFP28 ports labeled 48 through 51. Ports 0 through 47 support either 1-Gbps small form-factor pluggable (SFP) or 10-Gbps small form-factor pluggable plus (SFP+) transceivers. You can also use SFP+ DAC cables and 10-Gbps active optical cables (AOC) in any access port. The default 100-Gigabit Ethernet ports can be configured as 40-Gigabit Ethernet, and in this configuration can either operate as dedicated 40-Gigabit Ethernet ports or can be channelized to 4 independent 10-Gigabit Ethernet ports using copper or fiber breakout cables.
QFX5200-32C switch with Enhanced Layer 2 software	PIC 0 provides 32 QSFP28 ports. The 100-Gigabit Ethernet ports can be channelized to two 50-Gigabit Ethernet or four 25-Gigabit Ethernet ports. The default 100-Gigabit Ethernet ports can be configured as 40-Gigabit Ethernet and operate as 40-Gigabit Ethernet or be channelized to four 10-Gigabit Ethernet ports.
QFX10002-36Q switch with Enhanced Layer 2 software	PIC 0 provides 144 10-Gigabit Ethernet interfaces, and 36 40-Gbps QSFP+ ports, and 12 100-Gigabit Ethernet interfaces.
QFX10002-72Q switch with Enhanced Layer 2 software	PIC 0 provides 288 10-Gigabit Ethernet interfaces, and 72 40-Gbps QSFP+ ports, and 24 100-Gigabit Ethernet interfaces.
QFX10008 switch with Enhanced Layer 2 software	PIC 0 provides one-thousand, one-hundred fifty two 10-Gigabit Ethernet interfaces, two-hundred eighty-eight 40-Gbps QSFP+ ports, or two-hundred forty 100-Gigabit Ethernet interfaces.

Table 9: Naming Conventions for PICs (*continued*)

Device with Software Package	Convention
QFX10016 switch with Enhanced Layer 2 software	PIC 0 provides two-thousand, three-hundred and four 10-Gigabit Ethernet interfaces, five-hundred seventy-six 40-Gbps QSFP+ ports, or four-hundred eighty 100-Gigabit Ethernet interfaces.

- *port*—Interfaces use the following convention for port numbers:

Table 10: Naming Conventions for PORTs

Device with Software Package	Convention
QFX3500 switch with a QFabric software package	<p>There are 48 network access ports (10-Gigabit Ethernet) labeled 0 through 47 on PIC 0 and, 16 network access ports labeled 0 through 15 on PIC 1, and four 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 2. You can use the QSFP+ ports to connect the Node device to Interconnect devices.</p> <p>By default, the 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. Optionally, you can choose to configure the QSFP+ ports as 40-Gigabit Ethernet ports (see “Configuring the QSFP+ Port Type on QFX3500 Standalone Switches” on page 204).</p>
QFX3500 switch with Enhanced Layer 2 software	<p>There are 48 network access ports labeled 0 through 47 on PIC 0 and 4 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 1. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p>
QFX3600 switch with a QFabric software package	<p>There are 64 network access ports (10-Gigabit Ethernet) labeled Q0 through Q15 on PIC 0, and there are 16 network access ports (40-Gigabit Ethernet) labeled Q0 through Q15 on PIC 1.</p> <p>By default, all the QSFP+ ports are configured to operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the QSFP+ ports as 10-Gigabit Ethernet ports (see “Configuring the Port Type on QFX3600 Standalone Switches” on page 202) and use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.</p>

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
QFX3600 Node device with a QFabric software package	<p>PIC 0 can support up to 56 10-Gigabit Ethernet ports labeled Q2 through Q15, and PIC 1 can support up to 8 40-Gigabit data plane uplink ports labeled Q0 through Q7, and up to 14 40-Gigabit Ethernet ports labeled Q2 through Q15.</p> <p>On a QFX3600 Node device, by default, four 40-Gbps QSFP+ ports (labeled Q0 through Q3) are configured for uplink connections between your Node device and your Interconnect devices, and twelve 40-Gbps QSFP+ ports (labeled Q4 through Q15) use QSFP+ to four SFP+ copper breakout cables to support up to 48 10-Gigabit Ethernet ports for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and ports Q2 through Q15 for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks (see <i>Configuring the Port Type on QFX3600 Node Devices</i>).</p>
QFX3600 switch with Enhanced Layer 2 software	<p>PIC 0 can support 64 network access ports (10-Gigabit Ethernet ports) labeled Q0 through Q15 and 16 40-Gigabit Ethernet ports labeled Q0 through Q15. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p>
QFX5100-48S switch with Enhanced Layer 2 software	<p>PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47 and 6 40-Gbps QSFP+ ports labeled 48 through 53. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p>
EX4600 switch with Enhanced Layer 2 software	<p>PIC 0 can support 24 network access ports (10-Gigabit Ethernet ports) labeled 0 through 23 and 4 40-Gbps QSFP+ ports labeled 24 through 27. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provide 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 10-Gbps SFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p>

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
QFX5100-48S switch with a QFabric software package	<p>PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47, and PIC 1 can support 6 40-Gbps QSFP+ ports labeled 0 through 5. See “Configuring the QSFP+ Port Type on QFX5100 Devices” on page 207 for information on how to configure the port mode of 40-Gbps QSFP+ ports.</p>
QFX5100-24Q switch with Enhanced Layer 2 software	<p>PIC 0 can support 24 40-Gbps QSFP+ ports labeled 0 through 23. PIC 1 and PIC 2 each support 4 40-Gbps QSFP+ port, for a total of eight 40-Gbps QSFP+ ports. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>NOTE: You cannot channelize the 40-Gbps QSFP+ ports provided in the two QFX-EM-4Q expansion modules. Also, even though there is a total of 128 physical ports, only 104 logical ports can be channelized.</p> <p>You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See “Configuring the System Mode” on page 197 for information on how to configure the system mode.</p>
QFX5100-96S switch with Enhanced Layer 2 software	<p>PIC 0 can support 96 10-Gigabit Ethernet ports labeled 0 through 95, and 8 40-Gbps QSFP+ ports labeled 96 through 103. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>NOTE: You can only channelize the 40-Gbps QSFP+ ports provided in ports 96 and 100, because only 104 logical ports can be channelized.</p> <p>You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See “Configuring the System Mode” on page 197 for information on how to configure the system mode.</p>

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
QFX5110-48S switch with Enhanced Layer 2 software	<p>PIC 0 can support 48 10-Gigabit Ethernet ports labeled 0 through 47, and 4 QSFP28 ports labeled 48 through 51. These data ports (0 through 47) support either 1-Gbps small form-factor pluggable (SFP) or 10-Gbps small form-factor pluggable plus (SFP+) transceivers. You can also use SFP+ DAC cables and 10-Gbps active optical cables (AOC) in any access port. The default 100-Gigabit Ethernet ports can be configured as 40-Gigabit Ethernet, and in this configuration can either operate as dedicated 40-Gigabit Ethernet ports or can be channelized to 4 independent 10-Gigabit Ethernet ports using copper or fiber breakout cables.</p>
QFX5200-32C switch with Enhanced Layer 2 software	<p>There is support for both quad small-form-factor pluggable (QSFP+) and 28-Gbps QSFP+ (QSFP28) transceivers in the 32 QSFP28 sockets. The QSFP28 ports are configured as 100-Gigabit Ethernet ports by default, but can also be configured to speeds of 50, 40, 25, or 10 Gigabit Ethernet.</p> <p>The 100 Gigabit Ethernet ports can be channelized using breakout cables either to 2 independent downstream 50 Gigabit Ethernet or to 4 independent 25 Gigabit Ethernet ports. The default 100 Gigabit Ethernet ports can also be configured as 40 Gigabit Ethernet and in this configuration can either operate as dedicated 40 Gigabit Ethernet ports or can be channelized to 4 independent 10 Gigabit Ethernet ports using breakout cables. See “Channelizing Interfaces on QFX5200-32C Switches” on page 149 for information on how to configure and channelize the interfaces.</p> <p>NOTE: Autochannelization is not supported.</p>

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
QFX10002-36Q switch with Enhanced Layer 2 software	<p>There are 36 quad small-form factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 36 ports, 12 ports are QSFP28 capable, which are dual speed 40- or 100-Gigabit Ethernet optical transceivers.</p> <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using 28-Gbps QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. • 10-Gigabit Ethernet using breakout cables. When configured for channelization, a breakout cable converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports. <p>Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 12 QSFP28 ports support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers • 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports support:</p> <ul style="list-style-type: none"> • 40-Gigabit Ethernet QSFP+ transceivers • Access ports

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
QFX10002-72Q switch with Enhanced Layer 2 software	<p>There are 72 quad small-form factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 72 ports, 24 ports are QSFP28 capable, which are dual speed 40- or 100-Gigabit Ethernet optical transceivers.</p> <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using 28-Gbps QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. • 10-Gigabit Ethernet using breakout cables. When configured for channelization, a breakout cable converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports. <p>Any of the 72 ports 0 through 71 can be configured as either uplink or access ports. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 24 QSFP28 ports support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers <p>Each of the 72 QSFP+ ports support:</p> <ul style="list-style-type: none"> • 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports support:</p> <ul style="list-style-type: none"> • 40-Gigabit Ethernet QSFP+ transceivers • Access ports • Uplink ports

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
<p>On a QFX10008 switch with Enhanced Layer 2 software, there are two line cards available:</p> <p>QFX10008 with Line Card QFX10000-36Q (ELS)</p>	<p>QFX10000-36Q, a 36-port 40-Gigabit Ethernet quad small form-factor pluggable plus transceiver (QSFP+) or 12-port 100GbE QSFP28 line card</p> <p>The QFX10000-36Q line cards supports</p> <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. • 10-Gigabit Ethernet using breakout cabling and attached optical transceivers. When configured for channelization, the system converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports. <p>Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 12 QSFP28 ports supports:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers • 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 12 QSFP28 ports supports:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers • 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports support:</p> <ul style="list-style-type: none"> • 40-Gigabit Ethernet QSFP+ transceivers • Access ports • Uplink ports

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
QFX10008 with Line Card QFX10000-30C and QFX10000-30C-M (ELS)	<p>QFX10000-30C and QFX10000-30C-M, a 30-port 100-Gigabit or 40-Gigabit Ethernet QSFP28 line card</p> <ul style="list-style-type: none"> • The QFX10000-30C and QFX10000-30C-M line cards support: <p>Thirty 28-Gbps QSFP+ Pluggable Solution (QSFP28) cages that support either 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The QFX10000-30C and QFX10000-30C-M ports auto detect the type of transceiver installed and set the configuration to the appropriate speed.</p> <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. <p>See “Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134 for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 30 QSFP28 ports supports:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers • 40-Gigabit Ethernet QSFP+ transceivers • Access ports • Uplink ports

Table 10: Naming Conventions for PORTs *(continued)*

Device with Software Package	Convention
On a QFX10016 switch running Enhanced Layer 2 software, there are 16 slots, which you can populate with two types line cards: QFX10016 with Line Card QFX10000-36Q (ELS)	

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
	<ul style="list-style-type: none"> QFX10000-36Q, a 36-port 40-Gigabit Ethernet quad small form-factor pluggable plus transceiver (QSFP+) or 12-port 100GbE QSFP28 line card <p>The QFX10000-36Q line card consists of 36 quad small form-factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 36 ports, 12 ports are QSFP28 capable. The QSFP+ ports are dual speed and can support either 40-Gigabit or 100-Gigabit Ethernet optical transceivers. The line card can support 10-Gigabit Ethernet by channelizing the 40-Gigabit ports. Channelization is supported on fiber break-out cable using standard structured cabling techniques.</p> <p>With 100-Gigabit Ethernet using QSFP28 optical transceivers, when a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet.</p> <p>You can use 40-Gigabit Ethernet using QSFP+ optical transceivers.</p> <p>With 10-Gigabit Ethernet using breakout cabling and attached optical transceivers, when configured for channelization, the system converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports.</p> <p>Any of the 36 ports 0 through 35 can be configured as either uplink or access ports.</p> <p>Each of the 12 QSFP28 ports supports:</p> <ul style="list-style-type: none"> 100-Gigabit Ethernet QSFP28 transceivers 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports supports:</p> <ul style="list-style-type: none"> 40-Gigabit Ethernet QSFP+ transceivers Access ports <p>You can use 40-Gigabit Ethernet QSFP+ transceivers in any downstream port.</p> Uplink ports <p>You can configure all the QSFP+ ports as uplinks.</p> <p>Every second and sixth port in a 6XQSFP cage on a QFX10000-36Q supports 100-Gigabit Ethernet using QSFP28 transceivers. These 100-Gigabit Ethernet ports work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet, but are recognized as 40-Gigabit Ethernet by default. When a 40-Gigabit Ethernet transceiver is inserted into a 100-Gigabit Ethernet port, the port recognizes the 40-Gigabit Ethernet port speed. When a 100-Gigabit Ethernet transceiver is inserted into the port and enabled in the CLI, the port recognizes the 100-Gigabit Ethernet speed and disables two adjacent 40-Gigabit Ethernet ports. You can also use an 100-Gigabit Ethernet transceiver and run it at 40-Gigabit Ethernet by using the CLI to set the port speed to 40-Gigabit Ethernet.</p>

Table 10: Naming Conventions for PORTs (*continued*)

Device with Software Package	Convention
	<p>The 40-Gigabit Ethernet ports can operate independently, be channelized into four 10-Gigabit Ethernet ports, or bundled with the next two consecutive ports and channelized into twelve 10-Gigabit Ethernet ports as a port range. Only the first and fourth port in each 6XQSFP cage are available to channelize a port range. The port range must be configured using the <code>set chassis fpc 0 pic 0 port 1 channel-speed 10g</code> command. For example, to channelize the first switch port, use the set chassis fpc 0 pic 0 port 1 channel-speed 10g command.</p>
QFX10016 with Line Card QFX10000-30C and QFX10000-30C-M (ELS)	<p>The QFX10000-30C and QFX10000-30C-M line cards consist of thirty 28-Gbps QSFP+ Pluggable Solution (QSFP28) cages that support either 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The QFX10000-30C and QFX10000-30C-M ports auto detect the type of transceiver installed and set the configuration to the appropriate speed.</p> <p>Each QSFP28 socket supports:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into any of the ports, the QSFP28 socket is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. When a QSFP+ transceiver is inserted into any of the ports, the QSFP+ socket is enabled for 40-Gigabit. Any of the 30 ports 0 through 29 can be configured as either uplink or access ports, and of the 30 QSFP28 ports supports: <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers • 40-Gigabit Ethernet QSFP+ transceivers

Logical Part of an Interface Name on a Switch Running QFabric Software Package for QFX series, NFX Series, EX4600, QFabric System

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *device-name* (QFabric systems only): *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

Interface	State	VLAN members	Blocking
node-device1:xe-0/0/1.0	down	remote-analyzer	unblocked
node-device1:xe-0/0/2.0	down	default	unblocked
node-device1:xe-0/0/3.0	down	default	unblocked

When you configure aggregated Ethernet interfaces, you configure a logical interface, which is called a *bundle* or a *LAG*. Each LAG can include up to eight Ethernet interfaces, depending on the switch model.

Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software for QFX series, NFX Series, EX4600, QFabric System

Channelizing enables you to configure four 10-Gigabit Ethernet interfaces from a 40-Gigabit Ethernet QSFP+ interface. By default, a 40-Gigabit Ethernet QSFP+ interface is named *et-fpc/pic/port*. The resulting 10-Gigabit Ethernet interfaces appear in the following format: *xe-fpc/pic/port:channel*, where channel can be a value of 0 through 3.

For example, if an *et* interface named **et-0/0/3** is channelized to four 10-Gigabit Ethernet interfaces, the resulting 10-Gigabit Ethernet interface names will be **xe-0/0/3:0**, **xe-0/0/3:1**, **xe-0/0/3:2**, and **xe-0/0/3:3**:

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/3:0	up	down			
xe-0/0/3:1	up	down			
xe-0/0/3:2	up	down			
xe-0/0/3:3	up	down			

Wildcard Characters in Interface Names for QFX series, NFX Series, EX4600, QFabric System

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

Physical Part of an Interface Name for OCX1100

Interfaces in Junos OS are specified as follows:

type-fpc/pic/port

The convention is as follows:

- *type*—The OCX Series device interfaces use the following media types:
 - **xe**—10-Gigabit Ethernet interface
 - **et**—40-Gigabit Ethernet interface
 - **em**—Management interface
- *fpc*—Flexible PIC Concentrator. OCX Series interfaces use the following convention for the FPC number in interface names:

- On standalone OCX Series switches, the FPC number is always **0**.

The FPC number indicates the slot number of the line card that contains the physical interface.

- *pic*—The OCX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
 - PIC **0** provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces.
- *port*—Interfaces use the following convention for port numbers:
 - PIC **0** can support 48 network access ports (10-Gigabit Ethernet ports) labeled 1 through 48 and 6 40-Gbps QSFP+ ports labeled 49 through 54.

Wildcard Characters in Interface Names for OCX1100

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

SEE ALSO

[Interfaces Overview for Switches | 30](#)

[Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches | 134](#)

[Understanding Management Interfaces | 56](#)

[Understanding Port Ranges and System Modes | 167](#)

[Configuring the System Mode | 197](#)

[Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)

[Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support | 114](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

[Rear Panel of a QFX3500 Device](#)

[Front Panel of a QFX3600 Device](#)

Understanding Management Interfaces

You use management interfaces to access devices remotely. Typically, a management interface is not connected to the in-band network, but is connected to a device in the internal network. Through a management interface, you can access the device over the network using utilities such as **ssh** and **telnet**.

and configure it from anywhere, regardless of its physical location. As a security feature, users cannot log in as **root** through a management interface. To access the device as **root**, you must use the console port. You can also use **root** to log in using SSH.

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

Management interface port ranges vary based on device type (and platform support depends on the Junos OS release in your installation):

- QFX3500 devices:

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

- QFX3600 devices:

There are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0S** and **C1S**). On a QFX3600 standalone switch, the RJ-45 management interfaces and SFP management interfaces correspond to **me0** and **me1**. On a QFX3600 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**. Each pair of management interfaces correspond to one Ethernet interface—for example, both RJ-45 management interfaces (labeled **C0** and **C0S**) can correspond to **me0**, and both SFP management interfaces (labeled **C1** and **C1S**) can correspond to **me1**. By default, both RJ-45 management interfaces are active. If you insert an SFP interface into the SFP management port (**C0S**, for example), the SFP interface would become the active management interface, and the corresponding RJ-45 management interface (**C0**) is disabled.

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

- On QFX5100, QFX5200, and EX4600 switches, there is one RJ-45 management interface (labeled **C0**) and one SFP management interface (labeled **C1**), and they correspond to **em0** and **em1**. You can use both management interfaces simultaneously.
- On QFX10002 and QFX10008 switches, there is one RJ-45 management interface (labeled **MGMT**) and one SFP management interface (labeled **MGMT**), and they correspond to **em0** and **em1**. Although the CLI permits you to configure two management Ethernet interfaces within the same subnet, only one interface is usable and supported.

- On QFX10008 and QFX10016 switches, if you are using em1 for management purpose, then you cannot directly access the backup RE em1 from external network. Indirectly you can access the backup RE from external network, by following these steps:
 - Login to master RE using SSH/Telnet to its em1.
 - Access backup RE using the following command:

```
user@host>request routing-engine login other-routing-engine
```

- On OCX Series switches:

There is one RJ-45 management interface (labeled **MGMT**), which corresponds to em0. The em0 interface always has the status **up** in show command outputs, even if the physical port is empty. The me0 interface is a virtual interface between Junos and the host operating system, therefore its status is independent from the status of the physical port.

- QFabric system:

On a QFabric system, there are management interfaces on the Node devices, Interconnect devices, and Director devices. However, you cannot access the management interfaces on the Node devices or Interconnect devices directly. You can only manage and configure these devices using the Director device. You can connect to the management interface over the network using utilities such as SSH.

For information on how to use management interfaces on a QFabric system, see *Performing the QFabric System Initial Setup on a QFX3100 Director Group* and *Gaining Access to the QFabric System Through the Default Partition*.

Physical Interface Properties

IN THIS SECTION

- [Damping Shorter Physical Interface Transitions | 59](#)
- [Configuring Accounting for the Physical Interface | 60](#)
- [Enabling or Disabling SNMP Notifications on Physical Interfaces | 63](#)
- [Configuring Ethernet Loopback Capability | 64](#)
- [Configuring Short Reach Mode on QFX5100-48T | 65](#)
- [Configuring Flow Control | 66](#)
- [Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module | 67](#)

- [Setting the Operating Mode on a 2-Port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 Uplink Module | 68](#)
- [Configuring the Media Type on Dual-Purpose Uplink Ports | 71](#)
- [Disabling a Physical Interface | 72](#)

The physical interfaces undergo various transitions which is advertised to the Junos OS for proper functioning of the routers and switches. Accounting profiles that specify the characteristics of data about the traffic passing through the routers and switches can also be configured on the physical interfaces. Simple Network Management Protocol (SNMP) notifications can be enabled on the physical interface to provide information about the state of an interface or when a connection changes. The interface offers to configure various modes like short-reach-mode, flow-control and media type on the devices for ease of access.

Damping Shorter Physical Interface Transitions

By default, when an interface changes from being up to being down, or from down to up, this transition is advertised immediately to the hardware and Junos OS. In some situations—for example, when an interface is connected to an add/drop multiplexer (ADM) or wavelength-division multiplexer (WDM), or to protect against SONET/SDH framer holes—you might want to damp interface transitions. This means not advertising the interface's transition until a certain period of time has passed, called the *hold-time*. When you have damped interface transitions and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *down*, then the router begins to advertise the interface as being down. Similarly, when an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *up*, then the router begins to advertise the interface as being up. For information about physical interface damping, see *Physical Interface Damping Overview*.

This task applies to damping shorter physical interface transitions in milliseconds. To damp longer physical interface transitions in seconds, see *Damping Longer Physical Interface Transitions*.

To configure damping of shorter physical interface transitions:

1. Select the interface to damp, where the interface name is ***interface-type-fpc/pic/port***:

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

2. Configure the hold-time for link up and link down.

```
[edit interfaces interface-name]  
user@host# set hold-time up milliseconds down milliseconds
```

The hold time can be a value from 0 through 4,294,967,295 milliseconds. The default value is 0, which means that interface transitions are not damped. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

For most Ethernet interfaces, hold timers are implemented using a one-second polling algorithm. For 1-port, 2-port, and 4-port Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs), hold timers are interrupt-driven.

NOTE: The **hold-time** option is not available for controller interfaces.

SEE ALSO

SONET/SDH Defect Hold Times for Damping Interface Transitions Overview

Configuring SONET/SDH Defect Triggers

[hold-time](#) | [438](#)

Configuring Accounting for the Physical Interface

IN THIS SECTION

- [Accounting Profiles Overview](#) | [61](#)
- [Configuring Accounting for the Physical Interface](#) | [61](#)
- [Displaying Accounting Profile for the Physical Interface](#) | [62](#)

Accounting Profiles Overview

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management and Monitoring Guide*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter filter-name]** and **[edit firewall family family filter filter-name]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

Configuring Accounting for the Physical Interface

Before you begin

You must configure a profile to collect error and statistic information for input and output packets on a particular physical interface. An accounting profile specifies what statistics should be collected and written to a log file. For more information on how to configure an accounting-data log file, see the *Configuring Accounting-Data Log Files*.

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular physical interface.

1. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile profile-name]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set fieldsfield-name
```

2. Each accounting profile logs its statistics to a file in the **/var/log** directory. To configure which file to use, include the **file** statement at the **[edit accounting-options interface-profile profile-name]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set file filename
```

NOTE: You must specify a **file** statement for the interface profile that has already been configured at the **[edit accounting-options]** hierarchy level. For more information, see the [Configuring Accounting-Data Log Files](#)

3. Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the interval statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set interval minutes
```

NOTE: The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

4. To configure the interfaces on which the accounting needs to be performed, apply the interface profile to a physical interface by including the **accounting-profile** statement at the **[edit interfaces *interface-name*]** hierarchy level.

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

SEE ALSO

| [Configuring Accounting-Data Log Files](#)

Displaying Accounting Profile for the Physical Interface

Purpose

To display the configured accounting profile a particular physical interface at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level:

- interface-name—ge-1/0/1
- Interface profile —**if_profile**
- File name—**if_stats**
- Interval—15 minutes

Action

- Run the **show** command at the **[edit interfaces ge-1/0/1]** hierarchy level.

```
[edit interfaces ge-1/0/1]
accounting-profile if_profile;
```

- Run the **show** command at the **[edit accounting-options]** hierarchy level.

```
interface-profile if_profile {
  interval 15;
  file if_stats {
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-errors;
      output-errors;
    }
  }
}
```

Meaning

The configured accounting and its associated set options are displayed as expected.

Enabling or Disabling SNMP Notifications on Physical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. You can enable or disable these notification based on you requirements.

To explicitly enable sending SNMP notifications on the physical interface, perform the following steps:

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

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

2. Configure the **traps** option to enable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

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

To disable sending SNMP notifications on the physical interface, perform the following steps:

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

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

2. Configure the **no-traps** option to disable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

```
[edit interfaces interface-name]
user@host# set no-traps
```

SEE ALSO

| [traps](#) | [513](#)

Configuring Ethernet Loopback Capability

To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:


```
[edit]
user@switch# delete interfaces interface-name ether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

```
no-loopback;
```

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

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

SEE ALSO

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

Configuring Short Reach Mode on QFX5100-48T

You can enable short-reach mode for individual as well as a range of copper-based 10-Gigabit Ethernet interfaces using short cable lengths (less than 10m) on the QFX5100-48T switch. Short-reach mode reduces power consumption up to 5W on these interfaces.

1. To enable short-reach mode on an individual interface, issue the following command:

```
[edit chassis]
user@switch# set fpc fpc-slot pic pic-slot port port-number short-reach-mode enable
```

For example, to enable short-reach mode on port 0 on PIC 0, issue the following command:

```
[edit chassis]
user@switch# set fpc 0 pic 0 port 0 short-reach-mode enable
```

2. To enable short-reach mode on a range of interfaces, issue the following command:

```
[edit chassis]
user@switch# set fpc fpc-slot pic pic-slot port-range port-range-low port-range-high short-reach-mode enable
```

For example, to enable short-reach mode on a range of interfaces between port 0 and port 47 on PIC 0, issue the following command:

```
[edit chassis]
user@switch# set fpc 0 pic 0 port-range 0 47 short-reach-mode enable
```

3. To disable short-reach mode on an individual interface, issue the following command:

```
[edit chassis]
user@switch# set fpc fpc-slot pic pic-slot port port-number short-reach-mode disable
```

For example, to disable short-reach mode on port 0 on PIC 0, issue the following command:

```
[edit chassis]
user@switch# set fpc 0 pic 0 port 0 short-reach-mode disable
```

4. To disable short-reach mode on a range of interfaces, issue the following command:

```
[edit chassis]
user@switch# set fpc fpc-slot pic pic-slot port-range port-range-low port-range-high short-reach-mode disable
```

For example, to disable short-reach mode on a range of interfaces between port 0 and port 47 on PIC 0, issue the following command:

```
[edit chassis]
user@switch# set fpc 0 pic 0 port-range 0 47 short-reach-mode disable
```

SEE ALSO

| [short-reach-mode](#) | 660

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* gigheter-options]

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

SEE ALSO

flow-control

Ethernet Interfaces Overview

Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module

SFP+ uplink modules are supported on EX3200 and EX4200 switches, and SFP+ Media Access Control Security (MACSec) uplink modules are supported on EX4200 switches. You can use these uplink modules either for two SFP+ transceivers or four SFP transceivers. You configure the operating mode on the module to match the type of transceiver you want to use—that is, for SFP+ transceivers, you configure the 10-gigabit operating mode, and for SFP transceivers, you configure the 1-gigabit operating mode.

By default, the SFP+ uplink module operates in the 10-gigabit mode and supports only SFP+ transceivers. If you have not changed the module from the default setting and you want to use SFP+ transceivers, you do not need to configure the operating mode.

To set the operating mode of an SFP+ or SFP+ MACSec uplink module:

1. Change the operating mode to the appropriate mode for the transceiver type you want to use by using one of the following commands:

```
[edit]
```

```
user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 1g
```

```
[edit]
```

```
user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 10g
```

2. (SFP+ uplink module only) If the switch is running:

- Junos OS Release 10.1 or later, the changed operating mode takes effect immediately unless a port on the SFP+ uplink module is a Virtual Chassis port (VCP). If any port on the SFP+ uplink module is a VCP, the changed operating mode does not take effect until the next reboot of the switch.

NOTE: During the operating mode change, the Packet Forwarding Engine is restarted. In a Virtual Chassis configuration, this means that the Flexible PIC Concentrator connection with the master is dropped and then reconnected.

- Junos OS Release 10.0 or earlier, reboot the switch.

You can see whether the operating mode has been changed to the new mode you configured by issuing the **show chassis pic fpc-slot slot-number pic-slot 1** command.

SEE ALSO

Uplink Modules in EX3200 Switches

Uplink Modules in EX4200 Switches

Pluggable Transceivers Supported on EX3200 Switches

Pluggable Transceivers Supported on EX4200 Switches

Setting the Operating Mode on a 2-Port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 Uplink Module

You can configure the 2-port 4-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 uplink module on EX4300-48MP switches to operate either two 40-Gigabit Ethernet ports or two 100-Gigabit Ethernet port. By default, the uplink module operates only the two 40-Gbps ports.

The uplink module on EX4300-48MP switches supports Media Access Control Security (MACsec). See *Understanding Media Access Control Security (MACsec)* for more information.

The uplink module does not support configuring virtual chassis ports.

To set the operating mode on this uplink module:

1. Install the 2-port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 uplink module only in PIC slot 2 on the switch. Insert the uplink module in the chassis and check whether it is detected by issuing the **show chassis hardware** command.
2. Change the operating mode to 100-Gigabit Ethernet mode, by issuing the following command on the first port (port 0). The port then recognizes the 100-Gigabit speed and disables the adjacent 40-Gigabit Ethernet port. The adjacent 40-Gigabit Ethernet port is disabled only when port 0 is loaded with 100G optics.

```
[edit]
```

```
user@switch# set chassis fpc 0 pic 2 port 0 speed 100G
```

3. You can change the operating mode to 100-Gigabit Ethernet mode on the second (port 1) by using the following command. This command overrides the **set chassis fpc 0 pic 2 port 0 speed 100G** command to change the operating mode to 100-Gigabit Ethernet mode.

```
[edit]
```

```
user@switch# request chassis system-mode mode-2x100G
```

4. Optional: Check whether the operating mode has been changed to the new mode you configured by issuing the **show chassis pic fpc-slot 0 pic-slot 2** command.

NOTE: If you configure both the ports on the uplink module to operate at 100-Gbps speed, the four built-in QSFP+ ports on the switch are disabled.

Starting with Junos OS Release 19.1R1, in the 2-port 40-Gigabit Ethernet QSFP+/1-port 100-Gigabit Ethernet QSFP28 uplink module of EX4300-48MP switches, you can channelize the 100-Gigabit four independent 25-Gigabit Ethernet ports by using breakout cables. You can configure only port 0 of the uplink module as 25-Gigabit Ethernet port. Issue the command **set chassis fpc 0 pic 2 port 0 channel-speed 25g** to channelize the 100-Gigabit Ethernet uplink port to four 25-Gigabit Ethernet uplink ports.

Starting with Junos OS Release 19.3R1, you can configure the 2-port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 uplink module on EX4300-48MP switches to operate either two 40-Gigabit Ethernet ports or two 100-Gigabit Ethernet ports.

You can also channelize the 40-Gigabit Ethernet interfaces to four independent 10-Gigabit Ethernet interfaces using breakout cables. To channelize the 100-Gigabit Ethernet interfaces to operate as four independent 25-Gigabit Ethernet, specify the port number and channel speed

1. To configure the 100-Gigabit Ethernet uplink port to operate as a 25-Gigabit Ethernet interface, specify the port number and channel speed by using the following command:

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

For example, to configure port 0 to operate as a 25-Gigabit Ethernet port:

```
[edit chassis fpc 0 pic 2]  
user@switch# set port 0 channel-speed 25g
```

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

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

NOTE: If you configure both the ports on the uplink module to operate at 100-Gbps speed, the four QSFP+ ports on the switch are disabled.

SEE ALSO

| *Uplink Modules in EX4300 Switches*

Configuring the Media Type on Dual-Purpose Uplink Ports

EX2200-C switches and ACX1000 routers provide two dual-purpose uplink ports. Each dual uplink port is a single interface that offers a choice of two connections: an RJ-45 connection for a copper Ethernet cable and an SFP connection for a fiber-optic Ethernet cable. You can choose to use either connection, but only one connection can be active at a time.

By default, if you plug a transceiver into the SFP connector, the port becomes a fiber-optic Gigabit Ethernet port, even if a copper Ethernet cable is plugged into the RJ-45 connection as well. If a transceiver is not plugged into the SFP connector, the port defaults to a copper 10/100/1000 Ethernet port.

You can constrain the use of the port to one connection type by configuring the media type for the port to be either copper or fiber. When you configure a media type on the port, the port will no longer accept the alternate connection type. For example, if you configure the uplink port as a fiber port and then plug a copper Ethernet cable into the RJ-45 connector, the interface will not come up.

To configure the media type for an uplink port:

```
user@switch# set interfaces interface-name media-type (Dual-Purpose Uplink Ports) media-type
```

For example, to set the media type for uplink port **ge-0/1/0** to copper:

```
user@switch# set interfaces ge-0/1/0 media-type copper
```

NOTE: When you change the media type setting for a dual-purpose uplink port, it can take up to 6 seconds for the interface to appear in operational commands.

SEE ALSO

| *EX2200 Switches Hardware Overview*

Disabling a Physical Interface

IN THIS SECTION

- [Disabling a Physical Interface | 72](#)
- [Example: Disabling a Physical Interface | 73](#)
- [Effect of Disabling Interfaces on T series PICs | 74](#)

Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration.



CAUTION: Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.

To disable a physical interface:

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

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

2. Include the **disable** statement.

```
[edit interfaces at-fpc/pic/port ]
user@host# set disable
```


NOTE: On the router, when you use the **disable** statement at the **edit interfaces** hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet PICs with SFP and XFP transceivers do support it and the laser will be turned off when the interface is disabled.



WARNING: Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

Example: Disabling a Physical Interface

Sample interface configuration:

```
[edit interfaces]
user@host# show
ge-0/3/2 {
  unit 0 {
    description CE2-to-PE1;
    family inet {
      address 20.1.1.6/24;
    }
  }
}
```

Disabling the interface:

```
[edit interfaces ge-0/3/2]
user@host# set disable
```

Verifying the interface configuration:

```
[edit interfaces ge-0/3/2]
user@host# show
disable; # Interface is marked as disabled.
unit 0 {
  description CE2-to-PE1;
  family inet {
    address 20.1.1.6/24;
  }
}
```

```
}
```

Effect of Disabling Interfaces on T series PICs

The following table describes the effect of using the **set interfaces disable *interface_name*** statement on T series PICs.

Table 11: Effect of set interfaces disable <interface_name> on T series PICs

PIC Model Number	PIC Description	Type of PIC	Behaviour
PF-12XGE-SFPP	10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 Router)	5	Tx laser disabled
PF-24XGE-SFPP	10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (T4000 Router)	5	Tx laser disabled
PF-1CGE-CFP	100-Gigabit Ethernet PIC with CFP (T4000 Router)	5	Tx laser disabled
PD-4XGE-XFP	10-Gigabit Ethernet, 4-port LAN/WAN XFP	4	Tx laser disabled
PD-5-10XGE-SFPP	10-Gigabit LAN/WAN with SFP+	4	Tx laser disabled
PD-1XLE-CFP	40-Gigabit with CFP	4	Tx laser disabled
PD-1CE-CFP-FPC4	100-Gigabit with CFP	4	Tx laser disabled
PD-TUNNEL	40-Gigabit Tunnel Services	4	NA
PD-4OC192-SON-XFP	OC192/STM64, 4-port XFP	4	Tx laser not disabled
PD-1OC768-SON-SR	OC768c/STM256, 1-port	4	Tx laser not disabled

RELATED DOCUMENTATION

[disable](#) | 424

Logical Interface Properties

IN THIS SECTION

- [Configuring the Interface Address | 76](#)
- [Adding a Logical Unit Description to the Configuration | 78](#)
- [Configuring the Media MTU | 79](#)
- [Setting the Protocol MTU | 80](#)
- [Configuring the Interface Bandwidth | 81](#)
- [Enabling or Disabling SNMP Notifications on Logical Interfaces | 82](#)
- [Configuring Accounting for the Logical Interface | 82](#)
- [Disabling a Logical Interface | 85](#)

The logical interfaces can be configured and the description is displayed in the output of the **show** commands. Media maximum transmission unit (MTU) is automatically calculated when configuring an interface and can also be modified. Simple Network Management Protocol (SNMP) notifications can be enabled on the logical interface to provide information about the state of an interface or when a connection changes.

Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vpls** families, you never configure an address.

NOTE: The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, perform the following steps:

1. Configure the interface address at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level.

- To configure an IPv4 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.

You can also assign multiple IPv4 addresses on the same interface.

```
[edit interfaces ]
user@host# set interface-name unit logical-unit-number family inet address a.b.c.d/nn
```

NOTE:

- Juniper Networks routers and switches support **/31** destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports **/31** destination prefixes before configuration.
- You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point .
- By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.
- If you configure the same IP address on multiple interfaces in the same routing instance, Junos OS applies the configuration randomly on one of the interfaces. The other interfaces will remain without an IP address.

- To configure an IPv6 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet6 address *aaaa:bbbb:....:zzzz/nn*** statement at the **[edit interfaces]** hierarchy level.

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

NOTE:

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

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

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

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

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

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

4. [Optional] For interfaces that carry IP version 6 (IPv6) traffic, configure the host to assign itself a unique 64-Bit IP Version 6 interface identifier (EUI-64).

```
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family address
address]
user@host# set eui-64
```

Adding a Logical Unit Description to the Configuration

You can include a text description of each logical unit in the configuration file. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no impact on the interface's configuration. To add a text description, include the **description** statement:

```
description text;
```

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*]

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.

NOTE: You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See “Using DHCP Relay Agent Option 82 Information” in the *Junos OS Broadband Subscriber Management and Services Library*.

For information about describing physical interfaces, see *Configuring Interface Description*.

Configuring the Media MTU

The media maximum transmission unit (MTU) is the largest data unit that can be forwarded without fragmentation. The default media MTU size used on a physical interface depends on the encapsulation being used on that interface. For a listing of MTU sizes for each encapsulation type, see *Media MTU Sizes by Interface Type*.

To configure the media-MTU size:

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

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

2. Include the **mtu** statement.

```
[edit interfaces interface-name]
mtu bytes;
```

- If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. You configure the protocol MTU by including the **mtu** statement at the following hierarchy levels:
 - [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*]

NOTE:

- Changing the media MTU or protocol MTU causes an interface to be deleted and added again.
- Because tunnel services interfaces are considered logical interfaces, you cannot configure the MTU setting for the physical interface. This means you cannot include the `mtu` statement at the [edit interfaces *interface-name*] hierarchy level for the following interface types: generic routing encapsulation (gr-), IP-IP (ip-), loopback (lo-), link services (ls-), multilink services (ml-), and multicast (pe-, pd-). You can, however, configure the protocol MTU on all tunnel interfaces except virtual tunnel (vt) interfaces. Starting in Junos OS Release 17.1R3, you cannot configure the maximum transmission unit (MTU) size for vt interfaces because the **mtu bytes** option is deprecated for vt interfaces. Junos OS sets the MTU size for vt interfaces by default to unlimited.
- If you configure an MTU value by including the `mtu` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family mpls] hierarchy level, the configured value is used.

Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see *Media MTU Sizes by Interface Type*.

To modify the MTU for a particular protocol family, include the **mtu** statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- [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*]

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see *Encapsulation Overhead by Interface Encapsulation Type*. If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must

also reduce the protocol MTU size. (You configure the media MTU by including the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level.)

NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

SEE ALSO

| [Media MTU Overview](#)

Configuring the Interface Bandwidth

By default, the Junos OS uses the physical interface's speed for the MIB-II object, **ifSpeed**. You can configure the logical unit to populate the **ifSpeed** variable by configuring a bandwidth value for the logical interface. The **bandwidth** statement sets an informational-only parameter; you cannot adjust the actual bandwidth of an interface with this statement.

NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$$

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

To configure the bandwidth value for a logical interface, include the **bandwidth** statement:

```
bandwidth rate;
```

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*]

rate is the peak rate, in bps or cps. You can specify a value 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). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation **c**; values expressed in cells per second are converted to bits per second using the formula 1 cps = 384 bps. The value can be any positive integer. The **bandwidth** statement is valid for all logical interfaces, except multilink interfaces.

Enabling or Disabling SNMP Notifications on Logical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the logical interface, include the **traps** statement; to disable these notifications on the logical interface, include the **no-traps** statement:

```
(traps | no-traps);
```

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*]

Configuring Accounting for the Logical Interface

IN THIS SECTION

- [Accounting Profiles Overview | 83](#)
- [Configuring Accounting for the Logical Interface | 83](#)
- [Displaying Accounting Profile for the Logical Interface | 84](#)

Accounting Profiles Overview

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management and Monitoring Guide*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter filter-name]** and **[edit firewall family family filter filter-name]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

Configuring Accounting for the Logical Interface

Before you begin

You must configure a profile to collect error and statistic information for input and output packets on a particular logical interface. An accounting profile specifies what statistics should be collected and written to a log file. For more information on how to configure an accounting-data log file, see the *Configuring Accounting-Data Log Files*.

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular logical interface.

1. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile profile-name]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set fieldsfield-name
```

2. Each accounting profile logs its statistics to a file in the **/var/log** directory. To configure which file to use, include the **file** statement at the **[edit accounting-options interface-profile profile-name]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set file filename
```

NOTE: You must specify a **file** statement for the interface profile that has already been configured at the **[edit accounting-options]** hierarchy level. For more information, see the [Configuring Accounting-Data Log Files](#)

- Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the interval statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

```
[edit accounting-options interface-profile profile-name]
user@host# set interval minutes
```

NOTE: The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

- To configure the interfaces on which the accounting needs to be performed, apply the interface profile to a logical interface by including the **accounting-profile** statement at the **[edit interfaces interface-name unit *logical-unit-number*]** hierarchy level.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number accounting-profile profile-name
```

SEE ALSO

[Accounting Options Overview](#)

[Configuring Accounting-Data Log Files](#)

Displaying Accounting Profile for the Logical Interface

Purpose

To display the configured accounting profile a particular logical interface at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level:

- interface-name—ge-1/0/1
- Logical unit number—1
- Interface profile —**if_profile**
- File name—**if_stats**
- Interval—15 minutes

Action

- Run the **show** command at the **[edit interfaces ge-1/0/1 unit 1]** hierarchy level.

```
[edit interfaces ge-1/0/1 unit 1]
accounting-profile if_profile;
```

- Run the **show** command at the **[edit accounting-options]** hierarchy level.

```
interface-profile if_profile {
  interval 15;
  file if_stats {
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-errors;
      output-errors;
    }
  }
}
```

Meaning

The configured accounting and its associated set options are displayed as expected.

Disabling a Logical Interface

You can unconfigure a logical interface, effectively disabling that interface, without removing the logical interface configuration statements from the configuration. To do this, include the **disable** statement:

```
disable;
```

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*]

When an interface is disabled, a route (pointing to the reserved target “**REJECT**”) with the IP address of the interface and a 32-bit subnet mask is installed in the routing table. See *Routing Protocols*.

Example: Disabling a Logical Interface

Sample interface configuration:

```
[edit interfaces]
user@host# show
et-2/1/1 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    vlan-id 1000;
    family inet {
      address 11.0.0.20/24;
    }
  }
}
```

Disabling the interface:

```
[edit interfaces et-2/1/1 unit 0]
user@host# set disable
```

Verifying the interface configuration:

```
[edit interfaces et-2/1/1]
user@host# show
disable; # Interface is marked as disabled.
unit 0 {
  vlan-id 1000;
```

```
family inet {  
    address 11.0.0.20/24;  
}  
}
```

Interface Ranges

IN THIS SECTION

- [Understanding Interface Ranges for Switches | 87](#)
- [Configuring Interface Ranges for EX Series Switches with ELS | 91](#)

Interface ranges represent similar type of interfaces with common configurations that are grouped together. The ranges contain a name, a range and the configuration statements which are common to all the similar interfaces.

Understanding Interface Ranges for Switches

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

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

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

Starting in Junos OS Release 14.1X53-D15 and later, the common configurations defined in the interface range will not be overridden but appended to the local configuration. In Junos OS Releases prior to 14.1X53-D15, the common configurations defined in the interface range will be overridden by the local configuration.

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

[Table 12 on page 89](#) lists the configuration hierarchies for the EX Series, NFX, OCX, QFX Series, and QFabric Series.

Table 12: Configuration hierarchies for EX Series

Configuration Hierarchies for EX Series	Configuration Hierarchies for EX4600, NFX, QFX Series, and QFabric Systems	Configuration Hierarchies for EX Series with ELS
	<ul style="list-style-type: none"> • protocols isis interface • protocols sflow interfaces <p>NOTE: These statements are not supported on OCX Series switches.</p>	<ul style="list-style-type: none"> • forwarding-options analyzer <i>name</i> input egress interface • forwarding-options analyzer <i>name</i> input ingress interface • poe interface • protocols dot1x authenticator interface • protocols igmp interface • protocols isis interface • protocols layer2-control bpdu-block interface • protocols link-management peer <i>name</i> lmp-control-channel • protocols link-management te-link <i>name</i> interface • protocols lldp interface • protocols lldp-med interface • protocols mstp interface • protocols oam ethernet link-fault-management interface • protocols ospf area <i>area-id</i> interface • protocols pim interface • protocols router-advertisement interface • protocols router-discovery interface • protocols rsvp interface • protocols sflow interfaces • protocols vstp vlan <i>vlan-id</i> interface • switch-options redundant-trunk-group <i>group-name</i> interface • switch-options voip interface <p>For ELS details, see <i>Using the Enhanced Layer 2 Software CLI</i>.</p>

Table 12: Configuration hierarchies for EX Series (*continued*)

Configuration Hierarchies for EX Series	Configuration Hierarchies for EX4600, NFX, QFX Series, and QFabric Systems	Configuration Hierarchies for EX Series with ELS
<ul style="list-style-type: none"> • ethernet-switching-options analyzer <i>name</i> input egress interface • ethernet-switching-options analyzer <i>name</i> input ingress interface • ethernet-switching-options analyzer output interface • ethernet-switching-options bpdu-block interface • ethernet-switching-options interfaces • ethernet-switching-options redundant-trunk-group <i>group-name</i> interface • ethernet-switching-options secure-access-port interface • ethernet-switching-options voip interface • poe interface • protocols dot1x authentication interface • protocols gvrp interface • protocols igmp interface • protocols igmp-snooping vlan <i>vlan-name</i> interface • protocols isis interface • protocols link-management peer lmp-control-channel interface • protocols link-management te-link <i>name</i> interface • protocols lldp interface • protocols lldp-med interface • protocols mpls interface • protocols mstp interface • protocols mstp msti-<i>id</i> interface • protocols mstp msti-<i>id</i> vlan <i>vlan-id</i> interface • protocols oam ethernet link-fault-management interface • protocols ospf area • protocols pim interface 		

Table 12: Configuration hierarchies for EX Series (continued)

Configuration Hierarchies for EX Series	Configuration Hierarchies for EX4600, NFX, QFX Series, and QFabric Systems	Configuration Hierarchies for EX Series with ELS
<ul style="list-style-type: none">• protocols rip group <i>group-name</i> neighbor• protocols ripng group <i>group-name</i> neighbor• protocols router-advertisement interface• protocols router-discovery interface• protocols rsvp interface• protocols sflow interfaces• protocols stp interface• protocols vstp vlan <i>vlan-id</i> interface• vlans <i>vlan-name</i> interface		

SEE ALSO

<i>Configuring Interface Ranges</i>
<i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>
<i>Configuring Aggregated Ethernet Links (CLI Procedure)</i>
<i>Configuring a Layer 3 Subinterface (CLI Procedure)</i>
<i>interface-range</i>
Configuring Link Aggregation 281
<i>Configuring a Layer 3 Logical Interface</i>
<i>Junos OS Network Interfaces Library for Routing Devices</i>

Configuring Interface Ranges for EX Series Switches with ELS

NOTE: This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Interface Ranges*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

Junos OS allows you to group a range of identical interfaces into an *interface range*. You first specify the group of identical interfaces in the interface range. Then you can apply a common configuration to the specified interface range, reducing the number of configuration statements required and saving time while producing a compact configuration.

- [Configuring Interface Ranges on Switches | 92](#)
- [Expanding Interface Range Member and Member Range Statements | 95](#)
- [Configuration Inheritance for Member Interfaces | 96](#)
- [Member Interfaces Inheriting Configuration from Configuration Groups | 98](#)
- [Interfaces Inheriting Common Configuration | 99](#)
- [Configuring Inheritance Range Priorities | 100](#)
- [Configuration Expansion Where Interface Range Is Used | 101](#)

Configuring Interface Ranges on Switches

To configure an interface range, include the **interface-range** statement at the **[edit interfaces]** hierarchy level.

The **interface-range** statement accepts only physical networking interface names in its definition.

Interfaces can be grouped either as a range of interfaces or using a number range under the **interface-range** statement definition.

Interfaces in an **interface-range** definition can be added as part of a member range or as individual members or multiple members using a number range.

To specify a member range, use the **member-range** statement at the **[edit interfaces interface-range name]** hierarchy level.

To specify interfaces in lexical order, use the **member-range start-range to end-range** statement.

A range for a member statement must contain the following:

- *****—All, specifies sequential interfaces from 0 through 47.



CAUTION: The wildcard ***** in a member statement does not take into account the interface numbers supported by a specific interface type. Irrespective of the interface type, ***** includes interface numbers ranging from 0 through 47 to the interface group. Therefore, use ***** in a member statement with caution.

- **num**—Number; specifies one specific interface by its number.

- **[low-high]**—Numbers between low to high; specifies a range of sequential interfaces.
- **[num1, num2, num3]**—Numbers **num1**, **num2**, and **num3** specify multiple specific interfaces.

Example: Specifying an Interface Range Member Range

```
member-range ge-0/0/0 to ge-4/0/40;
```

To specify one or multiple members, use the **member** statement at the **[edit interfaces interface-range name]** hierarchy level.

To specify the list of interface range members individually or for multiple interfaces using regex, use the **member list of interface names** statement.

Example: Specifying an Interface Range Member

```
member ge-0/0/0;
member ge-0/*/*
member ge-0/[1-10]/0;
member ge-0/[1,2,3]/3;
```

Regex or wildcards are not supported for interface-type prefixes. For example, prefixes **ge**, **fe**, and **xe** must be mentioned explicitly.

An **interface-range** definition can contain both **member** and **member-range** statements within it. There is no maximum limit on the number of **member** or **member-range** statements within an interface-range. However, at least one **member** or **member-range** statement must exist within an **interface-range** definition.

Example: Interface Range Common Configuration

Configuration common to an interface range can be added as a part of the **interface-range** definition, as follows:

```
[edit]
interfaces {
  + interface-range foo {
  + member-range ge-1/0/0 to ge-4/0/40;
  + member ge-0/1/1;
  + member ge-5/[1-10]/*;
```

```

/*Common configuration is added as part of interface-range definition*/
mtu 256;
hold-time up 10;
ether-options {
    flow-control;
    speed {
        100m;
    }
    802.3ad primary;
}
}
}

```

An **interface-range** definition having just **member** or **member-range** statements and no common configurations statements is valid.

These defined interface ranges can be used in other configuration hierarchies, in places where an **interface** node exists.

Example: Interface-Range foo Used Under the Protocols Hierarchy

```

protocols {
    dot1x {
        authenticator {
            interface foo{
                retries 1;
            }
        }
    }
}

```

foo should be an **interface-range** defined at the **[interfaces]** hierarchy level. In the above example, the **interface** node can accept both individual interfaces and interface ranges.

TIP: To view an interface range in expanded configuration, use the **(show | display inheritance)** command. For more information, see the *CLI User Guide*.

The defined interface ranges can be used at places where the interface node is used. To view the configuration hierarchies, see [“Understanding Interface Ranges for Switches” on page 87](#).

Expanding Interface Range Member and Member Range Statements

All **member** and **member-range** statements in an interface range definition are expanded to generate the final list of interface names for the specified interface range.

Example: Expanding Interface Range Member and Member Range Statements

```
[edit]
interfaces {
  interface-range range-1 {
    member-range ge-0/0/0 to ge-4/0/20;
    member ge-10/1/1;
    member ge-5/[0-5]/*;
    /*Common configuration is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
}
```

For the **member-range** statement, all possible interfaces between **start-range** and **end-range** are considered in expanding the members. For example, the following **member-range** statement:

member-range ge-0/0/0 to ge-4/0/20

expands to:

```
[ge-0/0/0, ge-0/0/1 ... ge-0/0/max_ports
ge-0/1/0 ge-0/1/1 ... ge-0/1/max_ports
ge-0/2/0 ge-0/2/1 ... ge-0/2/max_ports
.
.
```

```

ge-0/MAX_PICS/0 ... ge-0/max_pics/max_ports
ge-1/0/0 ge-1/0/1 ... ge-1/0/max_ports
.
ge-1/MAX_PICS/0 ... ge-1/max_pics/max_ports
.
.
ge-4/0/0 ge-4/0/1 ... ge-4/0/max_ports]

```

The following **member** statement:

ge-5/[0-5]/*

expands to:

```

ge-5/0/0 ... ge-5/0/max_ports
ge-5/1/0 ... ge-5/0/max_ports
.
.
ge-5/5/0 ... ge-5/5/max_ports

```

The following **member** statement:

ge-5/1/[2,3,6,10]

expands to:

```

ge-5/1/2
ge-5/1/3
ge-5/1/6
ge-5/1/10

```

Configuration Inheritance for Member Interfaces

When the Junos OS expands the **member** and **member-range** statements present in an **interface-range**, it creates *interface objects* if they are not explicitly defined in the configuration. The common configuration is copied to all its member interfaces in the **interface-range**.

Example: Configuration Priorities

Foreground interface configuration takes priority compared to configuration inherited by the interface through the **interface-range**.

```

interfaces {
  interface-range range-1 {
    member-range ge-1/0/0/ to ge-10/0/47;
    mtu 256;
  }
  ge-1/0/1 {
    mtu 1024;
  }
}

```

In the preceding example, interface **ge-1/0/1** will have an MTU value of 1024.

This can be verified with output of the **show interfaces | display inheritance** command, as follows:

user@host: # **show interfaces | display inheritance**

```

## 'ge-1/0/0' was expanded from interface-range 'range-1'
##
ge-1/0/0 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
}
ge-1/0/1 {
  mtu 1024;
}
##
## 'ge-1/0/2' was expanded from interface-range 'range-1'
##
ge-1/0/2 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
}
.....
.....
##

```

```
## 'ge-10/0/47' was expanded from interface-range 'range-1'
##
ge-10/0/47 {
  ##
  ## '256' was expanded from interface-range 'range-1'
  ##
  mtu 256;
}
```

Member Interfaces Inheriting Configuration from Configuration Groups

Interface range member interfaces inherit the config-groups configuration like any other foreground configuration. **interface-range** is similar to any other foreground configuration statement. The only difference is that the **interface-range** goes through a member interfaces expansion before Junos OS reads this configuration.

```
groups {
  global {
    interfaces {
      <*> {
        hold-time up 10;
      }
    }
  }
  apply-groups [global];
  interfaces {
    interface-range range-1 {
      member-range ge-1/0/0 to ge-10/0/47;
      mtu 256;
    }
  }
}
```

The **hold-time** configuration is applied to all members of **interface-range range-1**.

This can be verified with **show interfaces | display inheritance** as follows:

```
user@host# show interfaces | display inheritance
```

```
ge-1/0/0 {
  ##
  ## '256' was expanded from interface-range 'range-1'
```

```

    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}
ge-1/0/1 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}
ge-10/0/47 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}

```

SEE ALSO

Using Wildcards with Configuration Groups

Interfaces Inheriting Common Configuration

If an interface is a member of several interface ranges, that interface will inherit the common configuration from all of those interface ranges.

[edit]

```

interfaces {
  interface-range range-1 {
    member-range ge-1/0/0 to ge-10/0/47;
    mtu 256;
  }
}
interfaces {
  interface-range range-1 {
    member-range ge-10/0/0 to ge-10/0/47;
    hold-time up 10;
  }
}

```

In this example, interfaces **ge-10/0/0** through **ge-10/0/47** will have both **hold-time** and **mtu**.

Configuring Inheritance Range Priorities

The interface ranges are defined in the order of inheritance priority, with the first interface range configuration data taking priority over subsequent interface ranges.

```

[edit]
interfaces {
  interface-range int-grp-one {
    member-range ge-0/0/0 to ge-4/0/40;
    member ge-1/1/1;
    /*Common config is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
  }
}
interfaces {
  interface-range int-grp-two {
    member-range ge-5/0/0 to ge-10/0/40;
    member ge-1/1/1;
    mtu 1024;
  }
}

```

Interface **ge-1/1/1** exists in both **interface-range int-grp-one** and **interface-range int-grp-two**. This interface inherits **mtu 256** from **interface-range int-grp-one** because it was defined first.

Configuration Expansion Where Interface Range Is Used

In this example, **interface-range range-1** is used under the **protocols** hierarchy:

```
[edit]
interfaces {
  interface-range range-1 {
    member ge-10/1/1;
    member ge-5/5/1;
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
}
protocols {
  dot1x {
    authenticator {
      interface range-1 {
        retries 1;
      }
    }
  }
}
}
```

The **interface** node present under **authenticator** is expanded into member interfaces of the **interface-range range-1** as follows:

```
protocols {
  dot1x {
    authenticator {
      interface ge-10/1/1 {
        retries 1;
      }
      interface ge-5/5/1 {
        retries 1;
      }
    }
  }
}
```

```
}
```

The **interface range-1** statement is expanded into two interfaces, ge-10/1/1 and ge-5/5/1, and configuration **retries 1** is copied under those two interfaces.

This configuration can be verified using the **show protocols dot1x | display inheritance** command.

RELATED DOCUMENTATION

| *Physical Interfaces*

Gigabit Ethernet Interface

SUMMARY

You can configure Gigabit Ethernet Interface with various modes like speed options, autonegotiation options, VLAN options, IP options, interface modes, link settings on the switches. The configuration uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style.

IN THIS SECTION

- [Speed and Autonegotiation | 102](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109](#)
- [Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support | 114](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches | 119](#)

Speed and Autonegotiation

IN THIS SECTION

- [Configure Interface Speed on Switches | 103](#)
- [Speed and Autonegotiation | 103](#)
- [Configure Speed and Autonegotiation on QFX5100-48T Switches | 105](#)

- Autonegotiation support for QFX5110-48S and QFX5100-48S with JNP-SFPP-10GE-T transceiver | 106
- Autonegotiation support for QFX5120-48Y with JNP-SFPP-10GE-T transceiver | 108

Configure Interface Speed on Switches

Following are the steps to configure the speeds on EX and QFX platforms:

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 speed;
```

Speed and Autonegotiation

[Table 13 on page 104](#) describes the autonegotiation speeds supported on switches.

Table 13: Autonegotiation support

Switch	Speed and Autonegotiation Information
EX2300-48MP and EX2300-24MP	<ul style="list-style-type: none"> • The mge interface is a rate-selectable (multirate) Gigabit Ethernet interface that 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. Note that 10 Mbps speed is supported only on ge interfaces of EX2300 switch. • You can configure both multi-rate gigabit ethernet interface (mge) and gigabit ethernet (ge) interface using the speed configuration statement. • The Multi-rate gigabit ethernet interface (MGE) on EX2300-24MP and EX2300-48MP switches flaps (becomes unavailable, and then available again) while performing timeout detection and recovery (TDR) test. • 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.
EX4300-48MP (EX-UM-4SFPP-MR)	<ul style="list-style-type: none"> • 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. • 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. • 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.

On 1/10G capable Gigabit Ethernet SFP interfaces, the duplex is always full and the speed matches that of the inserted optic. These interfaces support either 1G or 10G SFP optics.

- For EX and QFX products, the CLI configuration needs to match the optic speed. For 1G based optics, the configuration needs to start with **ge-**. For 10G inserted optics, the configuration needs to start with **xe-**. By default, both **ge** and **xe** choices are in the default configuration. User must match the CLI syntax to the optic speed.
- For MX products, the display and configuration is always **xe-** only, even if a 1G optic is inserted. The **xe-** value is used to denote that the interface is 10G capable. If a 1G optic is used, show commands for the interface will display the correct speed, but the configuration will always show as **xe-**.

NOTE: Only 10 Gbps and 40 Gbps interfaces are supported on OCX Series switches.

Configure Speed and Autonegotiation on QFX5100-48T Switches

For information about speed support, see [speed](#).

[Table 14 on page 105](#) provides QFX5100-48T details and description.

Table 14: QFX5100-48T details and description

Detail	Description
Duplex Mode	Full duplex
Autonegotiation	<p>The autonegotiation option is to negotiate the speeds.</p> <p>10 Gbps, 1 Gbps - By default, autonegotiation enabled</p> <p>100 Mbps - Autonegotiation not supported</p>

Following are guidelines for configuring speed and autonegotiation on QFX5100-48T switch:

- If the speed on the switch is set to 10-Gbps or auto, the switch advertises all the speeds.
- If the speed on the switch is set to 1-Gbps or 100-Gbps, the switch advertises only the respective speeds (1-Gbps or 100-Gbps).
- Whenever you connect a new device to a previously used interface, the switch operates in auto-negotiate mode, unless already set to a fixed speed.
- The **no-auto-negotiation** statement does no action. Hence, it is recommended not to use the **no-auto-negotiation** statement.

[Table 15 on page 106](#) provides the configuration steps to configure speed and autonegotiation.

Table 15: Configure Speed and Autonegotiation

Configure Speed/Autonegotiation	Use Configuration
To configure a particular speed, mention the speed.	<p>For a port to only advertise a specific speed, start with a specific speed, it is mandatory that both the auto-negotiation option must be set (enabled) and the interface must also be configured with a specific supported speed.</p> <pre>set interfaces xe-0/0/0 ether-options auto-negotiation set interfaces xe-0/0/0 speed <i>speed</i></pre> <p>For example to configure 1-Gbps speed, execute the following command:</p> <pre>set interfaces xe-0/0/0 ether-options auto-negotiation set interfaces xe-0/0/0 speed 1g</pre>
To enable auto-negotiation and advertise all speeds.	<p>With or without below, QFX5100-48T interface support auto-negotiation to one of either 10G and 1G.</p> <pre>set interfaces xe-0/0/0 ether-options auto-negotiation set interfaces xe-0/0/0 speed auto</pre> <p>This configuration does not change any functionality. If the speed is set to fixed at 10G, the interface still operates as auto, and advertises 10G/1G/100M.</p> <p>When you configure a port using the speed auto option, the port deletes the last configured speed, comes up again and advertises all the possible speeds.</p>

Autonegotiation support for QFX5110-48S and QFX5100-48S with JNP-SFPP-10GE-T transceiver

The interfaces on which the JNP-SFPP-10GE-T transceivers are present come up based on the speed (100-Mbps, or 1Gbps, or 10-Gbps) configured using the **set interfaces *interface-name* speed *speed*** command at the remote end.

For information about platforms support, see [Hardware Compatibility Tool](#).

[Table 16 on page 107](#) discusses QFX5110-48S and QFX5100-48S switches with JNP-SFPP-10GE-T transceiver details.

Table 16: QFX5110-48S and QFX5100-48S switches with JNP-SFPP-10GE-T transceiver details and description

Details	Description
QFX5110-48S and QFX5100-48S switches with JNP-SFPP-10GE-T transceiver	<p>Interface created - Multirate gigabit Ethernet (mge) interface.</p> <p>The mge interface is a rate-selectable (multirate) Gigabit Ethernet interface that can support 10-Gbps speed over CAT5e/CAT6/CAT6a cables. On QFX5110-48S and QFX5100-48S, the mge interface supports 10-Gbps, 1-Gbps, and 100-Mbps speeds, which can be configured using the speed configuration statement.</p> <p>Use the set interfaces mge-0/0/0 speed (100M 1G 10G) command to configure the speed and autonegotiation.</p>
QFX5110-48S and QFX5100-48S switches with other transceivers	Interface created - ge or the xe interface.
Duplex Mode	Full duplex
Viewing Media Specific Information	<p>You can execute the show interfaces media command to view the media-specific information. In the output of show interfaces name media, the output field speed displays the speed configured for the mge interface (with a default of 10G). The configured speed signifies the highest speed that the JNP-SFPP-10GE-T is capable of working at. You should enable auto-negotiation for the JNP-SFPP-10GE-T, unless it works in 100-Mbps speed where it can use the parallel detect capability using which it can detect when the link partner is in forced 100BASE-TX mode and bring the link up. The speed displayed under the Link partner denotes the actual speed at which the link is working. The Link partner speed is dynamic and displays the highest speed that both ends have negotiated and can work at.</p>

When the interface is configured with a particular speed, it means that the transceiver can support connection to a peer at rates lesser than or equal to the configured speed, as shown in [Table 17 on page 107](#):

Table 17: Interface speed based on configured speed and remote end speed

Configured Speed	The interface will be up with remote end speed of
10G	10G, 1G, and 100M
1G	1G and 100M
100M	100M

Autonegotiation support for QFX5120-48Y with JNP-SFPP-10GE-T transceiver

For information about platforms support, see [Hardware Compatibility Tool](#).

[Table 18 on page 108](#) discusses QFX5120-48Y with JNP-SFPP-10GE-T transceiver details.

Table 18: QFX5120-48Y with JNP-SFPP-10GE-T transceiver details and description

Details	Description
Supported speeds	<p>10 Gbps and 1 Gbps</p> <p>Default speed: 10 Gbps (with or without JNP-SFPP-10GE-T transceiver connected)</p> <p>If the peer does not support 10-Gbps speed, then the link will be down.</p>
Duplex Mode	Full duplex

[Table 19 on page 108](#) configure 1-Gbps and 10-Gbps speeds on QFX5120-48Y with JNP-SFPP-10GE-T transceiver.

Table 19: Configure and delete 1-Gbps speed

Configure Speed	Description
1 Gbps	<p>Use the set chassis fpc 0 pic 0 port <i>port-number</i> speed 1G command. Due to hardware limitations, you can configure the <i>port-number</i> value only in multiples of four, starting from port 0. You must also configure sets of four consecutive ports (for example, 0-3, 4-7, and so on) to operate at the common speed.</p> <p>On QFX5120 switch, multirate gigabit Ethernet (mge) interfaces are not supported due to hardware limitations.</p>
To revert to 10-Gbps speed (after setting 1-Gbps speed)	Delete the 1G speed configuration

SEE ALSO

[speed](#) | [505](#)

[auto-negotiation \(Switches\)](#) | [413](#)

[auto-negotiation \(Routers\)](#)

[Port Settings](#) | [130](#)

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Devices include a factory default configuration that:

- Enables all 10-Gigabit Ethernet network interfaces on the switch
- Sets a default port mode (access)
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching**
- Configures Storm Control on all 10-Gigabit Ethernet network interfaces
- Provides basic Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP) configuration

The **ether-options** statement enables you to modify the following options:

- **802.3ad**—Specify an aggregated Ethernet bundle for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.
- **autonegotiation**—Enable or disable autonegotiation of flow control, link mode, and speed for interfaces.
- **link-mode**—Specify **full-duplex**, **half-duplex**, or **automatic** for Gigabit Ethernet interfaces.
- **loopback**—Enable or disable a loopback interface for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.

To set **ether-options** for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces:

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

This topic describes:

- [Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 110](#)
- [Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches | 110](#)
- [Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches | 111](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches | 112](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches | 112](#)
- [Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 113](#)

Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches

If you are connecting a switch to other switches and to routers on the LAN, you need to assign the interface to a logical port and you need to configure the logical port as a trunk port.

To configure a Gigabit Ethernet or 10-Gigabit interface for trunk port mode on the Enhanced Layer 2 software (ELS):

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

SEE ALSO

| [Monitoring Interface Status and Traffic](#) | 392

Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches

Devices include a factory default configuration that enables Gigabit Ethernet interfaces with applicable link settings.

The following default configurations are available on Gigabit Ethernet interfaces:

- You cannot set the speed on these interfaces.
On QFX5100-48S and QFX5100-96S devices using 1-Gigabit Ethernet SFP interfaces, the speed is set to 1 Gbps by default and cannot be configured to operate in a different speed.
- On QFX5100 devices, the interface naming for Gigabit Ethernet interfaces changes automatically to xe-0/0/0, ge-0/0/0, or et-0/0/0 when the appropriate SFP is inserted.
- Gigabit Ethernet interfaces operate in full-duplex mode.
- Autonegotiation is supported by default. Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M and 1G. By default, autonegotiation is disabled on 10-Gigabit fiber ports.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.

NOTE: Do not use the **set interfaces *name* ether-options no-auto-negoatiate** command to remove the autonegotiation configuration.

Issue the **show interfaces *name* extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches

Devices include a factory default configuration that enables Gigabit Ethernet interfaces with applicable link settings.

The following default configurations are available on Gigabit Ethernet interfaces:

- Gigabit Ethernet interfaces operate in full-duplex mode.
- Gigabit Ethernet interfaces must be configured as *xe-fpc/pic/port*, and not *ge-fpc/pic/port*.
- Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 1G and 10G. By default, autonegotiation is enabled on 10-Gigabit Ethernet Copper interfaces.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command. 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: Do not use the **set interfaces *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

You can reenable autonegotiation it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

Issue the **show interfaces *name* extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

- 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.

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- All the 10-Gigabit Ethernet interfaces are set to **auto-negotiation**.
- Flow control for 10-Gigabit Ethernet interfaces is set to **enabled** by default. You can disable flow control by specifying the **no-flow-control** option.
- The speed cannot be configured.

On QFX5100-48S, QFX5100-96S, and QFX5100-24Q devices using 10-Gigabit Ethernet SFP interfaces, the speed is set to 10 Gbps by default and cannot be configured to operate in a different speed.

- On QFX5100 devices, the interface naming for Gigabit Ethernet interfaces changes automatically to xe-0/0/0, ge-0/0/0, or et-0/0/0 when the appropriate SFP is inserted.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M and 1G. By default, autonegotiation is disabled on 10-Gigabit fiber ports.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.

NOTE: Do not use the **set interfaces *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

Issue the **show interfaces *name* extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- All the 10-Gigabit Ethernet interfaces are set to **auto-negotiation**.
- Flow control for 10-Gigabit Ethernet interfaces is set to **enabled** by default. You can disable flow control by specifying the **no-flow-control** option.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 1G and 10G.

NOTE: On QFX5100-48T-6Q switches, autonegotiation is not supported for 100-Mbps speed.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

Issue the **show interfaces *name* extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches

To specify an IP address for the logical unit:

```
[edit]
```

```
user@switch# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

SEE ALSO

[Monitoring Interface Status and Traffic | 392](#)

[show interfaces xe | 1213](#)

[show interfaces ge | 1075](#)

[speed | 505](#)

Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support

NOTE: This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

An Ethernet interface must be configured for optimal performance in a high-traffic network. EX Series switches include a factory default configuration that:

- Enables all the network interfaces on the switch
- Sets a default interface mode (access)
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching** (except on EX8200 switches and Virtual Chassis)
- Specifies Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP)

This topic describes:

- [Configuring VLAN Options and Interface Mode | 114](#)
- [Configuring the Link Settings | 115](#)
- [Configuring the IP Options | 118](#)

Configuring VLAN Options and Interface Mode

By default, when you boot a switch and use the factory default configuration, or when you boot the switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that instead of **default**. You can also configure a port to accept untagged packets from the user-configured VLAN. For details on this concept (native VLAN), see *Understanding Bridging and VLANs on Switches*.

If you are connecting either a desktop phone, wireless access point or a security camera to a Power over Ethernet (PoE) port, you can configure some parameters for the PoE interface. PoE interfaces are enabled by default. For detailed information about PoE settings, see *Configuring PoE on EX Series Switches (CLI Procedure)*.

If you are connecting a device to other switches and to routers on the LAN, you need to assign the interface to a logical port and configure the logical port as a trunk port. See *Port Role Configuration with the J-Web Interface (with CLI References)* for more information about port configuration.

If you are connecting to a server that contains virtual machines and a VEPA for packet aggregation from those virtual machines, configure the port as a tagged-access port. See *Understanding Bridging and VLANs on Switches* for more information about tagged access.

To configure a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface for trunk port mode:

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

SEE ALSO

| *Monitoring Interface Status and Traffic*

Configuring the Link Settings

EX Series switches include a factory default configuration that enables interfaces with the link settings provided in [Table 20 on page 115](#).

Table 20: Factory Default Configuration Link Settings for EX Series Switches

Ethernet Interface	Autonegotiation	Flow Control	Link Mode	Link Speed
1 gigabit	Enabled	Enabled	Autonegotiation (full duplex or half duplex) For information about EX4300, see the Note below this table.	Autonegotiation (10 Mbps, 100 Mbps, or 1 Gbps)
10 gigabit (using a DAC cable)	Enabled	Enabled	Full duplex	10 Gbps
10 gigabit (using a fiber-optic cable)	Disabled	Enabled	Full duplex	10 Gbps
40 gigabit (using a DAC cable)	Enabled	Enabled	Full duplex	40 Gbps

Table 20: Factory Default Configuration Link Settings for EX Series Switches (*continued*)

Ethernet Interface	Autonegotiation	Flow Control	Link Mode	Link Speed
40 gigabit (using a fiber-optic cable)	Disabled	Enabled	Full duplex	40 Gbps

NOTE: On EX4300 switches, there is no **link-mode** configuration statement. The link-mode setting on an EX4300 switch is handled as follows:

- If the link partner is operating in half duplex, the EX4300 interface goes to half duplex.
- If the link partner is not capable of autonegotiation, the EX4300 interface goes to half duplex.
- If the link partner is capable of autonegotiation and is operating in full duplex, the EX4300 interface also works in full duplex.
- To force an EX4300 interface to stay in full-duplex mode, configure the interface's speed as 10 Mbps or 100 Mbps and also configure the interface with the **no-autonegotiation** statement.

To configure the link mode and speed settings for a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface:

NOTE: On EX4300 switches, there is no **link-mode** configuration statement. See information earlier in this document regarding how the link mode is set on EX4300 switches.

```
[edit]
```

```
user@switch# set interfaces interface-name
```

To configure additional link settings for a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface:

```
[edit]
```

```
user@switch# set interfaces interface-name ether-options
```

For detailed information about the FPC, PIC, and port numbers used for EX Series switches, see [“Understanding Interface Naming Conventions” on page 38](#).

Configurable link settings include:

- [802.3ad](#)—Specify an aggregated Ethernet bundle. See *Configuring Aggregated Ethernet Links (CLI Procedure)*.
- **auto-negotiation**—Enable or disable autonegotiation of flow control, link mode, and speed.

NOTE: Starting with Junos OS Releases 14.1X53-D40, 15.1R4, and 17.1R1, half-duplex communication is supported on all built-in network copper ports on EX4300 switches. *Half-duplex* is bidirectional communication; however, signals can flow in only one direction at a time. *Full-duplex* communication means that both ends of the communication can send and receive signals at the same time.

Half-duplex is configured by default on EX4300 switches. If the link partner is set to autonegotiate the link, then the link is autonegotiated to full duplex or half duplex. If the link is not set to autonegotiation, then the EX4300 link defaults to half duplex unless the interface is explicitly configured for full duplex.

To explicitly configure full duplex:

```
[edit]
user@switch# set interfaces interface-name speed 10m-or-100m
[edit]
user@switch# set interfaces interface-name ether-options no-auto-negotiation
```

To verify a half-duplex (or a full-duplex) setting:

```
user@switch> show interfaces interface-name extensive
```

- **flow-control**—Enable or disable flow control.
- **link-mode**—Specify full duplex, half duplex, or autonegotiation.

NOTE: On EX4300 switches, there is no **link-mode** configuration statement. See information earlier in this document regarding how the link mode is set on EX4300 switches.

- **loopback**—Enable or disable loopback mode.
- **speed**—Specify 10 Mbps, 100 Mbps, 1 Gbps, or autonegotiation.

Configuring the IP Options

To specify an IP address for the logical unit using IPv4:

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

To specify an IP address for the logical unit using IPv6:

[edit]

user@switch# **set interfaces** *interface-name* unit *logical-unit-number* family inet6 address *ip-address*

NOTE: Access interfaces on EX4300 switches are set to **family ethernet-switching** by default. You might have to delete this or any other user-configured family setting before changing the setting to **family inet** or **family inet6**.

SEE ALSO

[Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\)](#)

[Monitoring Interface Status and Traffic](#)

[show interfaces ge | 1075](#)

[show interfaces xe | 1213](#)

[Understanding Interface Naming Conventions | 38](#)

Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches

Devices include a factory default configuration that:

- Enables all 10-Gigabit Ethernet network interfaces on the switch
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching**
- Configures Storm Control on all 10-Gigabit Ethernet network interfaces

This topic describes:

- [Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces | 120](#)
- [Configuring the IP Options | 120](#)

Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces

Devices include a factory default configuration that enables 10-Gigabit Ethernet and interfaces with applicable link settings.

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- The speed for 10-Gigabit Ethernet interfaces is set to 10 Gbps by default. The speed cannot be configured.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is not supported.

The **ether-options** statement enables you to modify the following options:

- **802.3ad**—Specify an aggregated Ethernet bundle for 10-Gigabit Ethernet interfaces.
- **loopback**—Enable or disable a loopback interface for 10-Gigabit Ethernet interfaces.

To set **ether-options** for 10-Gigabit Ethernet interfaces:

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

Configuring the IP Options

To specify an IP address for the logical unit:

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

SEE ALSO

[Monitoring Interface Status and Traffic | 392](#)

[show interfaces xe | 1213](#)

[show interfaces ge | 1075](#)

[speed | 505](#)

Optical Transport Network (OTN) Interfaces

IN THIS SECTION

- [Understanding the QFX10K-12C-DWDM Line Card | 121](#)
- [Configuring OTN Interface Options on QFX10K-12C-DWDM | 124](#)

The QFX10K-12C-DWDM line card supports the optical transport interfaces (OTN) which is used for high end packet forwarding by cloud providers, service providers and enterprises. There are various optic-specific options that can be configured on the QFX10K-12C-DWDM line card including the forward error correction (FEC) mode and enabling the threshold crossing alarms.

Understanding the QFX10K-12C-DWDM Line Card

IN THIS SECTION

- [Software Features | 122](#)
- [OTN Alarms and Defects | 123](#)

The QFX10000-12C-DWDM 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 QFX10K-12C-DWDM line card is supported on Junos OS Release 17.2R1 and later.

The following sections explain the features of the QFX10K-12C-DWDM line card in detail:

Software Features

The following interface features are supported on the QFX10000-12C-DWDM:

- 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 QFX10K-12C-DWDM 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)

Configuring OTN Interface Options on QFX10K-12C-DWDM

The QFX10000-12C-DWDM 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 QFX10K-12C-DWDM line card is supported on Junos OS Release 17.2R1 and later.

Each QFX10K-12C-DWDM 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 634
<i>otn-options</i>
<i>signal-degrade</i>
<i>preemptive-fast-reroute</i>

Port Settings

SUMMARY

Learn about how to channelize speeds on different ports, configure port type, and configure system mode on switches.

IN THIS SECTION

- [Channelize Block of Ports or Individual Port | 131](#)
- [Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches | 134](#)
- [Channelizing Interfaces on QFX5110-48S Switches | 146](#)
- [Channelizing Interfaces on QFX5200-32C Switches | 149](#)
- [Channelizing Interfaces on QFX5210-64C Switches | 152](#)
- [Channelizing Interfaces on QFX5120-32C Switches | 154](#)
- [Channelizing Interfaces on QFX5120-48Y Switches | 159](#)
- [Channelizing Interfaces on QFX5120-48T Switches | 163](#)
- [Channelizing Interfaces on EX4650-48Y Switches | 165](#)
- [Understanding Port Ranges and System Modes | 167](#)
- [Configuring the System Mode | 197](#)
- [Configuring the Port Type on QFX3600 Standalone Switches | 202](#)
- [Configuring the QSFP+ Port Type on QFX3500 Standalone Switches | 204](#)
- [Configuring the QSFP+ Port Type on QFX5100 Devices | 207](#)

Channelize Block of Ports or Individual Port

You can either configure a port as a single interface (non channelized interface) or partition the port into smaller data channels or multiple interfaces (channelized interfaces). On switches platform, you can channelize an individual port or block of ports or quad ports based on the platform.

[Table 21 on page 131](#) describes the steps to follow in channelizing a port and setting the port speed.

Table 21: Channelize block of ports or individual port

Configuration Steps	Details	Example
Channelize Individual Port: Configure an individual port to operate at a specific channel speed. Specify a port number and channel speed.	[edit chassis fpc <i>fpc-slot</i> pic <i>pic-slot</i>] user@switch# set port <i>port-number</i> channel-speed <i>speed</i>	To configure an individual 40-Gigabit Ethernet (et) port to operate as 10-Gigabit Ethernet (xe) ports, specify a port number and channel speed [edit chassis fpc 0 pic 0] user@switch# set port 3 channel-speed 10g
Channelize Block of Ports: Channelize a block of ports. Specify a port range and channel speed.	[edit chassis fpc <i>fpc-slot</i> pic <i>pic-slot</i>] user@switch# set port-range <i>port-range-low port-range-high</i> channel-speed <i>speed</i>	To configure ports 0 through 3 on PIC 0 to operate as 50-Gigabit Ethernet ports: [edit chassis fpc 0 pic 0] user@switch# set port-range 0 3 channel-speed 50g
Configure Speed per Quad: Configure port speeds only per quad (group of 4 ports) and not individually. Specify the speed for the first port of the quad ports. All ports operate at a single speed within the quad.	[edit chassis fpc <i>fpc-slot</i> pic <i>pic-slot</i>] user@switch# set port <i>port-number</i> speed <i>speed</i>	To configure ports 4 through 7 to operate as 25-Gigabit Ethernet ports, you must configure port 4 to operate as 25-Gigabit Ethernet ports. [edit chassis fpc 0 pic 0] user@switch# set port 4 speed 25g

Table 21: Channelize block of ports or individual port (*continued*)

Configuration Steps	Details	Example
(Optional) Return to default configuration, for an individual port or quad ports.	[edit chassis fpc <i>fpc-slot</i> pic <i>pic-slot</i>] user@switch# delete <i>port-number</i> <i>speed speed</i>	[edit chassis fpc 0 pic 0] user@switch# delete 4 speed 25G
(Optional) Return to default configuration, for block of ports.	[edit chassis fpc <i>fpc-slot</i> pic <i>pic-slot</i>] user@switch# delete port-range <i>port-range-low port-range-high</i> <i>channel-speed speed</i>	To return ports 0 through 3 from the 50-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration: [edit chassis fpc 0 pic 0] user@switch# delete port-range 0 3 channel-speed 50g
Commit the configuration.		

Interface Naming Conventions

Each interface name includes a unique identifier and follows a naming convention. When you configure the interface, use the interface name. You can either configure a port as a single interface (non channelized interface) or partition the port into smaller data channels or multiple interfaces (channelized interfaces).

When multiple interfaces are supported on a physical port, you use the colon (:) notation in the interface naming conventions as a delimiter to differentiate the multiple interfaces on a physical port. In the interface naming convention, **xe-x/y/z:channel**:

- x refers to the FPC slot number.
- y refers to the PIC slot number.
- z refers to the physical port number.
- channel refers to the number of channelized interfaces.

When the 40-Gigabit Ethernet interfaces (**et-fpc/pic/port**) are channelized as 10-Gigabit Ethernet interfaces, the interface appears in the **xe-fpc/pic/port:channel** format, and channel is a value of 0 through 3.

[Table 22 on page 133](#) describes the naming formats for the channelized and non-channelized interfaces.

Table 22: Channelized and Non-Channelized Interface Naming Formats

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

SEE ALSO

channel-speed 649
fpc 650
pic 655

Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches

NOTE: On QFX10008 and QFX10016 switches, channelization is supported on fiber break-out cables using standard structured cabling techniques.

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

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

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

NOTE: If you are using a EX4600 or QFX5100 switch with a peer QFX10000 switch, you must disable auto-channelization on the EX4600 or QFX5100 switch to ensure that the relevant interfaces are created and the links are up.

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

There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches, depending on the system mode you enable. If you try to channelize ports that are restricted, the configuration is ignored. See [“Configuring the System Mode” on page 197](#) for more information.

On QFX10002, QFX10008, and QFX10016 switches, there are 100-Gigabit Ethernet ports that work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet, but are recognized as 40-Gigabit Ethernet by default. You cannot channelize the 100-Gigabit Ethernet ports when they are operating as 100-Gigabit Ethernet interfaces. The 40-Gigabit Ethernet ports can operate independently or be channelized into four

10-Gigabit Ethernet ports as part of a port range. Ports cannot be channelized individually. Only the first and fourth port in each 6XQSFP cage is available to channelize as part of a port range. In a port range, the ports are bundled with the next two consecutive ports. For example, if you want to channelize ports 0 through 2, you would channelize port 0 only. If you try to channelize a port that is not supported, you will receive an error message when you commit the configuration. Auto-channelization is not supported on any ports.

When a 40-Gigabit Ethernet transceiver is inserted into a 100-Gigabit Ethernet port, the port recognizes the 40-Gigabit Ethernet port speed. When a 100-Gigabit Ethernet transceiver is inserted into the port and enabled in the CLI, the port recognizes the 100-Gigabit Ethernet speed and disables two adjacent 40-Gigabit Ethernet ports.

[Table 23 on page 136](#) provides detailed information on which ports are 100-Gigabit Ethernet, which ports can be channelized, and which ports are disabled when a 100-Gigabit Ethernet is inserted in the QFX10002-36Q switch and the QFX10000-36Q line card on a QFX10008 or QFX10016 switch.

[Table 24 on page 138](#) provides detailed information on which ports are 100-Gigabit Ethernet, which ports can be channelized, and which ports are disabled when a 100-Gigabit Ethernet is inserted in the QFX10002-72Q switch. On the QFX10008 and QFX10016 switches with the QFX10000-36Q line card installed, only ports 0 through 35 are available. For more information, see *QFX10002-72Q Port Panel* and *QFX10000-36Q Line Card*.

Table 23: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
0	✓	✓	✓	-	-
1	✓		✓	✓	0, 2
2	✓		✓	-	-
3	✓	✓	✓	-	-
4	✓		✓	-	-
5	✓		✓	✓	3, 4
6	✓	✓	✓	-	-
7	✓		✓	✓	6, 8
8	✓		✓	-	-

Table 23: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings (*continued*)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
9	✓	✓	✓	-	-
10	✓		✓	-	-
11	✓		✓	✓	9, 10
12	✓	✓	✓	-	-
13	✓		✓	✓	12, 14
14	✓		✓	-	-
15	✓	✓	✓	-	-
16	✓		✓	-	-
17	✓		✓	✓	15, 16
18	✓	✓	✓	-	-
19	✓		✓	✓	18, 20
20	✓		✓	-	-
21	✓	✓	✓	-	-
22	✓		✓	-	-
23	✓		✓	✓	21, 22
24	✓	✓	✓	-	-
25	✓		✓	✓	24, 26
26	✓		✓	-	-

Table 23: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings (*continued*)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
27	✓	✓	✓	-	-
28	✓		✓	-	-
29	✓		✓	✓	27, 28
30	✓	✓	✓	-	-
31	✓		✓	✓	30, 32
32	✓		✓	-	-
33	✓	✓	✓	-	-
34	✓		✓	-	-
35	✓		✓	✓	33, 34

Table 24: QFX10002-72Q Switch Port Mappings

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
0	✓	✓	✓	-	-
1	✓		✓	✓	0, 2
2	✓		✓	-	-
3	✓	✓	✓	-	-
4	✓		✓	-	-
5	✓		✓	✓	3, 4

Table 24: QFX10002-72Q Switch Port Mappings (continued)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
6	✓	✓	✓	-	-
7	✓		✓	✓	6, 8
8	✓		✓	-	-
9	✓	✓	✓	-	-
10	✓		✓	-	-
11	✓		✓	✓	9, 10
12	✓	✓	✓	-	-
13	✓		✓	✓	12, 14
14	✓		✓	-	-
15	✓	✓	✓	-	-
16	✓		✓	-	-
17	✓		✓	✓	15, 16
18	✓	✓	✓	-	-
19	✓		✓	✓	18, 20
20	✓		✓	-	-
21	✓	✓	✓	-	-
22	✓		✓	-	-
23	✓		✓	✓	21, 22

Table 24: QFX10002-72Q Switch Port Mappings (continued)

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

Table 24: QFX10002-72Q Switch Port Mappings (continued)

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

Table 24: QFX10002-72Q Switch Port Mappings (continued)

Port Number	4X10 Gigabit Ethernet Port	4X10 Gigabit Channelized Port Group	40-Gigabit Ethernet (Default)	100-Gigabit Ethernet	100-Gigabit Ethernet Disables
60	✓	✓	✓	-	-
61	✓		✓	✓	60, 62
62	✓		✓	-	-
63	✓	✓	✓	-	-
64	✓		✓	-	-
65	✓		✓	✓	63, 64
66	✓	✓	✓	-	-
67	✓		✓	✓	66, 68
68	✓		✓	-	-
69	✓	✓	✓	-	-
70	✓		✓	-	-
71	✓		✓	✓	69, 70

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

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

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

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

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

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

NOTE: The **port-range** statement is not available on QFX10002-72Q, QFX10008, and QFX10016 switches. Instead, configure the port range using the **port** statement. Starting from port 0, you channelize every third port to channelize a group of three ports. For example, channelize port 0 to channelize ports 0 through 2, port 3 to channelize ports 3 through 5, and so on. See [Table 23 on page 136](#) for port mapping information.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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


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

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

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

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

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

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

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

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

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

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

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

SEE ALSO

[Configuring the System Mode | 197](#)

[channel-speed | 649](#)

[fpc | 650](#)

[pic | 655](#)

Channelizing Interfaces on QFX5110-48S Switches

On the QFX5110-48S switch, there are four ports labeled 48 through 51, which support QFSP28 ports. The QSFP28 ports support 100-Gigabit Ethernet interfaces and 40-Gigabit Ethernet interfaces. You can channelize the 40-Gigabit Ethernet interfaces to four independent 10-Gigabit Ethernet interfaces using breakout cables.

NOTE: You cannot configure channelized interfaces to operate as Virtual Chassis ports.

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

NOTE: On QFX5110-48S standalone switches, the FPC value is always 0.

Starting in Junos OS release 20.1R1, in addition to 1-Gbps, 10-Gbps, 40-Gbps, 100-Gbps speeds, now you can also configure 100-Mbps speed using the **`set interfaces interface-name speed 100M`** command. By default, all 48 ports on QFX5110-48S come up with 10-Gbps speed. With QFX-SFP-1GE-T connected, along with 1-Gbps speed, now you can also configure 100-Mbps on QFX5110-48S switches.

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

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

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

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

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

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

```
[edit]
```

```
user@switch# commit
commit complete
```

3. To return a range of ports from the 10-Gigabit Ethernet configuration to the 40-Gigabit Ethernet configuration, delete the 10g statement:

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

For example, to return ports 48 through 51 from the 10-Gigabit Ethernet configuration to the 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port-range 48 51 channel-speed 10g
```

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

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

5. To return an individual 10-Gigabit Ethernet port to the 40-Gigabit Ethernet configuration, delete the 10g statement:

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

For example, to return port 48 from the 10-Gigabit Ethernet configuration to the 40-Gigabit Ethernet configuration:

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

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

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

SEE ALSO

[channel-speed](#) | [649](#)[fpc](#) | [650](#)[pic](#) | [655](#)

Channelizing Interfaces on QFX5200-32C Switches

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

NOTE: You cannot configure channelized interfaces to operate as Virtual Chassis ports.

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

By default, the 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces appear in the `et-fpc/pic/port` format. When the 100-Gigabit Ethernet interfaces are channelized as 50-Gigabit Ethernet and 25-Gigabit Ethernet interfaces, the interface names appear in the `et-fpc/pic/port:channel` format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the `xe-fpc/pic/port:channel` format, where channel can be a value of 0 through 3. To channelize the ports, manually configure the port speed using the **`set chassis fpc slot-number port port-number channel-speed speed`** command, where the speed can be set to 10G, 25G, or 50G. If a 100-Gigabit Ethernet transceiver is connected, you can only set the speed to 25G or 50G. If a 40-Gigabit Ethernet transceiver is connected, you can only set the speed to 10g. There is no commit check for this, however.

On the QFX5200-32C switches, the ports support auto-channelization starting in Junos OS Release 15.1X53-D230.

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

NOTE: On QFX5200-32C standalone switches, the FPC value is always 0.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

SEE ALSO

[channel-speed](#) | [649](#)

[fpc](#) | [650](#)

[pic](#) | [655](#)

Channelizing Interfaces on QFX5210-64C Switches

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

On the QFX5210-64C switch, there are a total of 64 physical ports. Any port can be used as either 100-Gigabit Ethernet or 40-Gigabit Ethernet interfaces. You choose the speed by plugging in the appropriate transceiver. QFX5210-64C switches can also be channelized to 50G, 25G or 10G. QFX5210-64C switch supports three port speeds in total, 10-Gbps, 40-Gbps, and 100-Gbps. It supports native 40, 100-Gigabit Ethernet configuration and channelized 10, 25, 50-Gigabit Ethernet configuration.

Starting in Junos OS Release 19.1R1, you can channelize the 100-Gigabit Ethernet interfaces to four independent 25-Gigabit Ethernet interfaces or two independent 50-Gigabit Ethernet interfaces. Based on the optics used, the 100-Gigabit Ethernet or 40-Gigabit Ethernet interface is supported. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces. This configuration can either operate as dedicated 40-Gigabit Ethernet interfaces, or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables on the QFX5210-64C switch. You can use all the available ports on QFX5210-64C switch without disabling the ports. The port channelization on QFX5210 switches occurs automatically when the total number of ports does not exceed 128 BCM ports and when the number of port per pipe does not exceed 32 BCM ports.

By default, the 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces are displayed in the `et-fpc/pic/port` format. When the 100-Gigabit Ethernet interfaces are channelized as 50-Gigabit Ethernet and 25-Gigabit Ethernet interfaces, the interface names appear in the `et-fpc/pic/port:channel` format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the `xe-fpc/pic/port:channel` format.

To channelize the ports, manually configure the port speed using the **`set chassis fpc slot-number port port-number channel-speed speed`** command, where the speed can be set to **10G**, or **25G**, or **50G**. If a 100-Gigabit Ethernet transceiver is connected, you can only set the speed to **25G** or **50G**. If a 40-Gigabit Ethernet transceiver is connected, you can only set the speed to **10G**. There is no commit check for this, however.

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

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

1. Configure a block of 100-Gigabit Ethernet (*et*) ports to operate as 50-Gigabit Ethernet ports, by specifying

the port range and channel speed:

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

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

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

2. Configure a block of 100-Gigabit Ethernet (et) ports to operate as 25-Gigabit Ethernet ports, by specifying the port range and channel speed:

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

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

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

3. Configure an individual 40-Gigabit Ethernet (et) port to operate as 10-Gigabit Ethernet (xe) ports, by specifying the port number and channel speed:

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

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

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

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

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

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

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

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

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

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

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

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

SEE ALSO

[channel-speed](#) | [649](#)

[fpc](#) | [650](#)

[pic](#) | [655](#)

Channelizing Interfaces on QFX5120-32C Switches

For information on QFX5120-32C Switches, see *QFX5120 Switch Hardware Guide*.

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

[Table 25 on page 155](#) QFX5120-32C details and description.

Table 25: QFX5120-32C details and description

Details	Description
FPC/PIC	FPC 0 and PIC 0; one FPC and one PIC.
QSFP/QSFP28 and SFP+ ports	Total number of ports - 34; 32 QSFP/QSFP28 ports and two 10G SFP+ ports.
Supported Port Speeds	<ul style="list-style-type: none"> 40-Gbps and 100-Gbps (non-channelized mode) 10-Gbps, 25-Gbps, and 50-Gbps (channelized mode)
Auto speed detection mode (Enabled by default)	<ul style="list-style-type: none"> Auto speed detection mode detects the 100-Gigabit Ethernet interface, 40-Gigabit Ethernet interface and channelize based on the configuration (auto channelization is enabled by default). If you have disabled auto-channelization, then to channelize the ports, manually configure the port speed using the set chassis fpc slot-number port port-number channel-speed speed command, where the speed can be set to 10G, or 25G, or 50G.

[Table 26 on page 155](#) summarizes the supported port speeds on QFX5120-32C.

Table 26: Port Speed for QFX5120-32C

PIC	Port Number	Port Speed Supported	Default Speed
PIC 0	0-30	40-Gigabit Ethernet 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet 2x50-Gigabit Ethernet 100-Gigabit Ethernet Refer Table 21 on page 131 to configure the channel speeds.	All active ports operate in 100 Gbps mode. You cannot configure channelized interfaces to operate as Virtual Chassis ports.
	Port 31	40-Gigabit Ethernet 2x50-Gigabit Ethernet 100-Gigabit Ethernet (Not Supported) 4x10-Gigabit Ethernet or 4x25-Gigabit Ethernet	

To channelize block of ports or individual ports on QFX5120-32C switches, see [Table 21 on page 131](#).

[Table 27 on page 157](#) describes the naming formats for the channelized and non-channelized interfaces on QFX5120-32C switches for the QSFP/QSFP28 ports (port 0 - 31).

Table 27: Interface Naming Convention for the QFX5120-32C Switch

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe-0/0/0:[0-3]	et-0/0/0:[0-3]	et-0/0/0	et-0/0/0:[0-1]	et-0/0/0
	xe-0/0/1:[0-3]	et-0/0/1:[0-3]	et-0/0/1	et-0/0/1:[0-1]	et-0/0/1
	xe-0/0/2:[0-3]	et-0/0/2: [0-3]	et-0/0/2	et-0/0/2: [0-1]	et-0/0/2
	xe-0/0/3:[0-3]	et-0/0/3: [0-3]	et-0/0/3	et-0/0/3: [0-1]	et-0/0/3
	xe-0/0/4:[0-3]	et-0/0/4: [0-3]	et-0/0/4	et-0/0/4: [0-1]	et-0/0/4
	xe-0/0/5:[0-3]	et-0/0/5: [0-3]	et-0/0/5	et-0/0/5: [0-1]	et-0/0/5
	xe-0/0/5:[0-3]	et-0/0/6: [0-3]	et-0/0/6	et-0/0/6: [0-1]	et-0/0/6
	xe-0/0/7:[0-3]	et-0/0/7: [0-3]	et-0/0/7	et-0/0/7: [0-1]	et-0/0/7
	xe-0/0/8:[0-3]	et-0/0/8: [0-3]	et-0/0/8	et-0/0/8: [0-1]	et-0/0/8
	xe-0/0/9:[0-3]	et-0/0/9: [0-3]	et-0/0/9	et-0/0/9: [0-1]	et-0/0/9
	xe-0/0/10:[0-3]	et-0/0/10: [0-3]	et-0/0/10	et-0/0/10: [0-1]	et-0/0/10
	xe-0/0/11:[0-3]	et-0/0/11: [0-3]	et-0/0/11	et-0/0/11: [0-1]	et-0/0/11
	xe-0/0/12:[0-3]	et-0/0/12: [0-3]	et-0/0/12	et-0/0/12: [0-1]	et-0/0/12
	xe-0/0/13:[0-3]	et-0/0/13: [0-3]	et-0/0/13	et-0/0/13: [0-1]	et-0/0/13
	xe-0/0/14:[0-3]	et-0/0/14: [0-3]	et-0/0/14	et-0/0/14: [0-1]	et-0/0/14
	xe-0/0/15:[0-3]	et-0/0/15: [0-3]	et-0/0/15	et-0/0/15: [0-1]	et-0/0/15
	xe-0/0/16:[0-3]	et-0/0/16: [0-3]	et-0/0/16	et-0/0/16: [0-1]	et-0/0/16
	xe-0/0/17:[0-3]	et-0/0/17: [0-3]	et-0/0/17	et-0/0/17: [0-1]	et-0/0/17
	xe-0/0/18:[0-3]	et-0/0/18: [0-3]	et-0/0/18	et-0/0/18: [0-1]	et-0/0/18
	xe-0/0/19:[0-3]	et-0/0/19: [0-3]	et-0/0/19	et-0/0/19: [0-1]	et-0/0/19

Table 27: Interface Naming Convention for the QFX5120-32C Switch (*continued*)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
	xe-0/0/20:[0-3]	et-0/0/20: [0-3]	et-0/0/20	et-0/0/20: [0-1]	et-0/0/20
	xe-0/0/21:[0-3]	et-0/0/21: [0-3]	et-0/0/21	et-0/0/21: [0-1]	et-0/0/21
	xe-0/0/22:[0-3]	et-0/0/22: [0-3]	et-0/0/22	et-0/0/22: [0-1]	et-0/0/22
	xe-0/0/23:[0-3]	et-0/0/23: [0-3]	et-0/0/23	et-0/0/23: [0-1]	et-0/0/23
	xe-0/0/24:[0-3]	et-0/0/24: [0-3]	et-0/0/24	et-0/0/24: [0-1]	et-0/0/24
	xe-0/0/25:[0-3]	et-0/0/25: [0-3]	et-0/0/25	et-0/0/25: [0-1]	et-0/0/25
	xe-0/0/25:[0-3]	et-0/0/26: [0-3]	et-0/0/26	et-0/0/26: [0-1]	et-0/0/26
	xe-0/0/27:[0-3]	et-0/0/27: [0-3]	et-0/0/27	et-0/0/27: [0-1]	et-0/0/27
	xe-0/0/28:[0-3]	et-0/0/28: [0-3]	et-0/0/28	et-0/0/28: [0-1]	et-0/0/28
	xe-0/0/29:[0-3]	et-0/0/29: [0-3]	et-0/0/29	et-0/0/29: [0-1]	et-0/0/29
	xe-0/0/30:[0-3]	et-0/0/30: [0-3]	et-0/0/30	et-0/0/30: [0-1]	et-0/0/30
	xe-0/0/31:[0-3]	et-0/0/31: [0-3]	et-0/0/31	et-0/0/31: [0-1]	et-0/0/31

See [“Interface Naming Conventions” on page 132](#) for channelized and non-channelized interface naming formats.

SEE ALSO

[channel-speed](#) | [649](#)

[fpc](#) | [650](#)

[pic](#) | [655](#)

Channelizing Interfaces on QFX5120-48Y Switches

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

NOTE: You cannot configure channelized interfaces to operate as Virtual Chassis ports.

On the QFX5120 switch, there are a total of 56 ports. Of these 56 ports, 8 ports (labeled 48 through 55) are uplink ports that support 100-Gigabit Ethernet interfaces (QSFP28 ports) and 40-Gigabit Ethernet interfaces (QSFP+ ports). The other 48 ports (labeled 0 through 47) are SFP+ ports that support 25-Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, or 1-Gigabit Ethernet interfaces. The default speed for the SFP+ ports is 10 Gbps.

By default, the 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces appear in the `et-fpc/pic/port` format. When the 100-Gigabit Ethernet interfaces are channelized as 25-Gigabit Ethernet interfaces, the interface names appear in the `et-fpc/pic/port:channel` format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the `xe-fpc/pic/port:channel` format, where channel can be a value of 0 through 3. Auto-channelization is enabled by default on the uplink ports. If you have disabled auto-channelization, then to channelize the ports, manually configure the port speed using the **`set chassis fpc slot-number port port-number channel-speed speed`** command, where the speed can be set to 10G or 25G. If a 100-Gigabit Ethernet transceiver is connected, you can only set the speed to 25G. For the SFP+ ports, you can set the speed to 25G or 1G. There is no commit check for this, however.

NOTE: On QFX5120 switches, the uplink ports support auto-channelization.

Starting in Junos OS Release 19.4R1, you can use any of the following JNP-SFP-25G-DAC cables to set 10-Gbps speed on the SFP28 ports of a QFX5120-48Y switch:

- JNP-SFP-25G-DAC-1M
- JNP-SFP-25G-DAC-3M
- JNP-SFP-25G-DAC-5M

If you've plugged a JNP-SFP-25G-DAC cable into a QFX5120-48Y switch, then the SFP28 ports come up with 10-Gbps speed by default. To configure the SFP28 ports to operate at 25-Gbps speed, you must explicitly configure the speed of the first port in the port group using the **`set chassis fpc 0 pic 0 port port-num speed 25g`** command.

The following steps describe how to channelize uplink ports (block of ports or individual ports).

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

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

For example, to configure ports 48 through 56 on PIC 0 to operate as 25-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]  
user@switch# set port-range 48 56 channel-speed 25g
```

2. To configure a block of 40-Gigabit Ethernet (et) ports to operate as 10-Gigabit Ethernet ports, specify a port range and channel speed:

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

For example, to configure ports 48 through 56 on PIC 0 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]  
user@switch# set port-range 48 56 channel-speed 10g
```

3. To configure an individual 100-Gigabit Ethernet (et) port to operate as 25-Gigabit Ethernet (xe) ports, specify a port number and channel speed:

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

For example, to configure port 55 to operate as a 25-Gigabit Ethernet port:

```
[edit chassis fpc 0 pic 0]  
user@switch# set port 55 channel-speed 25g
```

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

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


For example, to configure port 55 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]  
user@switch# set port 55 channel-speed 10g
```

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

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

6. To return a range of ports from the 25-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration, delete the 25g statement:

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

For example, to return ports 48 through 56 from the 25-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]  
user@switch# delete port-range 48 56 channel-speed 25g
```

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

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

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

The following steps describe how to channelize SFP/SFP28 block of ports. You can configure the SFP and SFP28 port speeds only per quad (group of 4 ports) and not individually. All ports operate at a single speed within the quad. For instance, if you configure ports 8 through 11 to operate as 25-Gigabit Ethernet ports and you insert a 10G SFP+ transceiver in port 10, an interface is not created for this port.

1. To configure a block of four 10-Gigabit Ethernet (*et*) ports (quad ports) to operate as 25-Gigabit Ethernet ports, specify the speed for the first port of the quad ports. For instance,

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

For example, to configure ports 4 through 7 to operate as 25-Gigabit Ethernet ports, you must configure port 4 to operate as 25-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port 4 speed 25g
```

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

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

3. To return a range of ports from the 25-Gigabit Ethernet configuration to the default 10-Gigabit Ethernet configuration, delete the 25g statement:

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

For example, to return ports 4 through 7 from the 25-Gigabit Ethernet configuration to the default 10-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port 4 speed 25g
```

NOTE: To configure the ports to another speed, you must delete the current port speed statement to return to the default 10-Gigabit Ethernet configuration.

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

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

SEE ALSO

[fpc | 650](#)

[pic | 655](#)

Channelizing Interfaces on QFX5120-48T Switches

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

On the QFX5120-48T switch, there a total of 54 ports. [Table 28 on page 163](#) summarizes the speed and autonegotiation support on QFX5120-48T switches.

Table 28: Speed and channelization support on QFX5120-48T switches

Ports	Speed Supported	Default Speed	AutoChannelization/Channelization using CLI
Port 0 - 47	1 Gbps and 10 Gbps	10 Gbps	Not Supported
Port 48 - 53	100 Gbps and 40 Gbps	100 Gbps	<p>Supports both manual (using CLI) channelization and auto-channelization. However, the manual/CLI channelization always takes the higher precedence.</p> <p>You can channelize a particular port, using the set chassis fpc 0 pic 0 port <i>port-num</i> channel-speed (10G 25G 50G) command.</p> <p>The port-range CLI is not supported, which means, you cannot channelize block of ports. You can only channelize the individual ports.</p> <p>All 40-Gigabit and 100-Gigabit Ethernet ports support auto channelization with the following limitations:</p> <ul style="list-style-type: none"> • Port 50 and 51 supports either 4x10G or 4x25G based on the optic used. • All ports support 2x50G.

To channelize individual ports on QFX5120-48T Switches, see [Table 21 on page 131](#).

[Table 29 on page 164](#) lists the supported channelization speeds on the QSFP ports.

Table 29: Supported channelization speeds on QSFP Ports

QSFP Port Number	2X50G	4X10G	4X25G	40G	100G
48	Supported	Not Supported	Not Supported	Supported	Supported
49	Supported	Not Supported	Not Supported	Supported	Supported
50	Supported	Supported	Supported	Supported	Supported
51	Supported	Supported	Supported	Supported	Supported
51	Supported	Not Supported	Not Supported	Supported	Supported
53	Supported	Not Supported	Not Supported	Supported	Supported

[Table 30 on page 164](#) describes the naming formats for the channelized and non-channelized interfaces on QFX5120-48T Switches.

Table 30: Interface Naming Convention for QSFP port on QFX5120-48T Switch

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	50-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	Not supported	Not supported	et-0/0/48	et-0/0/48:[0-1]	et-0/0/48
	Not supported	Not supported	et-0/0/49	et-0/0/49:[0-1]	et-0/0/49
	xe-0/0/50:[0-3]	et-0/0/50:[0-3]	et-0/0/50	et-0/0/50:[0-1]	et-0/0/50
	xe-0/0/51:[0-3]	et-0/0/51:[0-3]	et-0/0/51	et-0/0/51:[0-1]	et-0/0/51
	Not supported	Not supported	et-0/0/52	et-0/0/52:[0-1]	et-0/0/52
	Not supported	Not supported	et-0/0/53	et-0/0/53:[0-1]	et-0/0/53

See [“Interface Naming Conventions” on page 132](#) for channelized and non-channelized interface naming formats.

SEE ALSO

channel-speed 649
fpc 650
pic 655

Channelizing Interfaces on EX4650-48Y Switches

For information on EX4650-48Y Switches, see *EX4650 Switch Hardware Guide*.

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

[Table 31 on page 165](#) EX4650-48Y details and description.

Table 31: EX4650-48Y details and description

Details	Description
FPC/PIC	FPC 0 and PIC 0; one FPC and one PIC.
QSFP/QSFP28 and SFP+ ports	Total number of ports - 56; 48 SFP+ ports and eight uplink ports.
Auto speed detection mode (Enabled by default)	<p>If you have disabled auto-channelization, then to channelize the ports, manually configure the port speed using the set chassis fpc slot-number port port-number channel-speed speed command, where the speed can be set to 10G or 25G. If a 100-Gigabit Ethernet transceiver is connected, you can only set the speed to 25G. For the SFP+ ports, you can set the speed to 25G or 1G. There is no commit check for this, however.</p> <p>On EX4650 switches, the uplink ports support auto-channelization.</p>

[Table 32 on page 166](#) summarizes the supported port speeds on EX4650-48Y.

Table 32: Port Speed for EX4650-48Y

PIC	Port Number	Port Speed Supported	Default Speed
PIC 0	(labeled 0 through 47) 48 SFP+ ports	1-Gigabit Ethernet 10-Gigabit Ethernet 25-Gigabit Ethernet You can configure the SFP and SFP28 port speeds only per quad (group of 4 ports) and not individually. To configure speed for the quad, refer Table 21 on page 131 . The interface will not get created automatically on inserting 1-Gigabit Ethernet or 25-Gigabit Ethernet transceivers. You must use the CLI to configure the port speed to 1-Gigabit Ethernet or 25-Gigabit Ethernet mode manually.	10 Gbps
	(labeled 48 through 55) 8 uplink ports	100-Gigabit Ethernet (QSFP28 ports) 40-Gigabit Ethernet (QSFP+ ports) 4x10-Gigabit Ethernet 4x25-Gigabit Ethernet	100 Gbps (for QSFP28 ports) 40 Gbps (for QSFP+ ports)

To channelize blocks of ports or individual ports on uplink ports, see [Table 21 on page 131](#). To configure speed per quad on the SFP+ ports, [Table 21 on page 131](#).

[Table 33 on page 166](#) lists the interface naming conventions of SFP+ ports (labeled 0 through 47) for the EX4650-48Y switch.

Table 33: Interface Naming Convention for the EX4650-48Y Switch (SFP+ ports)

PIC	1-Gigabit Ethernet Interface	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface
0	ge-0/0/0	xe-0/0/0	et-0/0/0
	ge-0/0/1	xe-0/0/1	et-0/0/1
	ge-0/0/2	xe-0/0/2	et-0/0/2
	ge-0/0/3	xe-0/0/3	et-0/0/3
	ge-0/0/4	xe-0/0/4	et-0/0/4

Table 34 on page 167 lists the interface naming conventions of uplink ports (labeled 48 through 55) for the EX4650-48Y switch.

Table 34: Interface Naming Convention for the EX4650-48Y Switch (Uplink ports)

PIC	10-Gigabit Ethernet Interface	25-Gigabit Ethernet Interface	40-Gigabit Ethernet Interface	100-Gigabit Ethernet Interface
0	xe-0/0/48:[0-3]	et-0/0/48:[0-3]	et-0/0/48	et-0/0/48
	xe-0/0/49:[0-3]	et-0/0/49:[0-3]	et-0/0/49	et-0/0/49
	xe-0/0/50:[0-3]	et-0/0/50:[0-3]	et-0/0/50	et-0/0/50
	xe-0/0/51:[0-3]	et-0/0/51:[0-3]	et-0/0/51	et-0/0/51
	xe-0/0/52:[0-3]	et-0/0/52:[0-3]	et-0/0/52	et-0/0/52
	xe-0/0/53:[0-3]	et-0/0/53:[0-3]	et-0/0/53	et-0/0/53
	xe-0/0/54:[0-3]	et-0/0/54:[0-3]	et-0/0/54	et-0/0/54
	xe-0/0/55:[0-3]	et-0/0/55:[0-3]	et-0/0/55	et-0/0/55

See “Interface Naming Conventions” on page 132 for channelized and non-channelized interface naming formats.

SEE ALSO

channel-speed		649
fpc		650
pic		655

Understanding Port Ranges and System Modes

IN THIS SECTION

- Port Ranges for Different Media Types | 168
- Supported System Modes | 195

QFX Series devices and EX4600 switches can support different port ranges depending on the device, media type of the interface, the software that is running on the device, and the system mode.

This topic describes:

Port Ranges for Different Media Types

The following media types support the following port ranges:

- On a QFX3500 device:
 - The valid port range for a Fibre Channel (fc) interface is **0** through **5** and **42** through **47** on PIC **0**, with a total of 12 available Fibre Channel ports.

NOTE: Fibre Channel ports are not supported on QFX3500, QFX3600, and QFX5100 switches running Enhanced Layer 2 software.

- The valid port range for a Gigabit Ethernet (ge) interface is **6** through **41** on PIC **0** because the ports between **0** and **5** and **42** and **47** are reserved as Fibre Channel ports. The total number of available Gigabit Ethernet ports is 36, because 12 of the remaining 48 ports are reserved for Fibre Channel and 10-Gigabit Ethernet interfaces. Fibre Channel ports cannot be configured as Gigabit Ethernet ports.
- The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **47** on PIC **0**. The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **15** on PIC **1**. The total number of available 10-Gigabit Ethernet ports is 64.
- The valid port range for a 40-Gigabit data plane uplink interface is **0** through **3** on PIC **1**
- The valid port range for a 40-Gigabit Ethernet interface is **0** through **3** on PIC **2**. There are four available ports.
- On a QFX3600 Node device:
 - The valid port range for a 10-Gigabit Ethernet interface is **8** through **63** on PIC **0**. There are 56 available ports.
 - The valid port range for a 40-Gigabit Ethernet interface is **2** through **15** on PIC **1**. There are 14 available ports.
 - The valid port range for a 40-Gigabit data plane uplink interface is **0** through **7** on PIC **1**. There are eight available ports.

See [Table 37 on page 178](#) for physical port to logical port mappings.

- On a QFX3600 switch running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **63** on PIC **0**. There are 64 available ports.

- The valid port range for a 40-Gigabit Ethernet interface is **0** through **15** on PIC **0**. There are 16 available ports.

See [Table 38 on page 181](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC **0**. There are 48 available ports. When you channelize the 6 40-Gbps QSFP+ ports on **0** through **5** on PIC **1**, there are 72 available ports.

NOTE: On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

- The valid port range for a 40-Gbps QSFP+ port is **0** through **5** on PIC **1**. There are six available ports.

See [Table 40 on page 187](#) for physical port to logical port mappings.

- On EX4600 switches running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **23** on PIC **0**. There are 24 available ports. When you channelize the 4 40-Gbps QSFP+ ports on **24** through **27** on PIC **0**. There are 40 available ports.

See [Table 40 on page 187](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running a QFabric software package:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC **0**. There are 48 available ports.
 - The valid port range for a 40-Gbps QSFP+ port is **0** through **5** on PIC **1**. There are six available ports.

NOTE: On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

See [Table 41 on page 191](#) for physical port to logical port mappings.

- For QFX5100-24Q and QFX5100-96S switches running Enhanced Layer 2 Software, see [Table 42 on page 196](#) for physical port to logical port mappings for different system modes.

Table 35: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
0	fc-0/0/0	Not supported on this port	xe-0/0/0	Not supported on this port	Not supported on this port
1	fc-0/0/1	Not supported on this port	xe-0/0/1	Not supported on this port	Not supported on this port
2	fc-0/0/2	Not supported on this port	xe-0/0/2	Not supported on this port	Not supported on this port
3	fc-0/0/3	Not supported on this port	xe-0/0/3	Not supported on this port	Not supported on this port
4	fc-0/0/4	Not supported on this port	xe-0/0/4	Not supported on this port	Not supported on this port
5	fc-0/0/5	Not supported on this port	xe-0/0/5	Not supported on this port	Not supported on this port
6	Not supported on this port	ge-0/0/6	xe-0/0/6	Not supported on this port	Not supported on this port
7	Not supported on this port	ge-0/0/7	xe-0/0/7	Not supported on this port	Not supported on this port
8	Not supported on this port	ge-0/0/8	xe-0/0/8	Not supported on this port	Not supported on this port
9	Not supported on this port	ge-0/0/9	xe-0/0/9	Not supported on this port	Not supported on this port
10	Not supported on this port	ge-0/0/10	xe-0/0/10	Not supported on this port	Not supported on this port
11	Not supported on this port	ge-0/0/11	xe-0/0/11	Not supported on this port	Not supported on this port
12	Not supported on this port	ge-0/0/12	xe-0/0/12	Not supported on this port	Not supported on this port

Table 35: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
13	Not supported on this port	ge-0/0/13	xe-0/0/13	Not supported on this port	Not supported on this port
14	Not supported on this port	ge-0/0/14	xe-0/0/14	Not supported on this port	Not supported on this port
15	Not supported on this port	ge-0/0/15	xe-0/0/15	Not supported on this port	Not supported on this port
16	Not supported on this port	ge-0/0/16	xe-0/0/16	Not supported on this port	Not supported on this port
17	Not supported on this port	ge-0/0/17	xe-0/0/17	Not supported on this port	Not supported on this port
18	Not supported on this port	ge-0/0/18	xe-0/0/18	Not supported on this port	Not supported on this port
19	Not supported on this port	ge-0/0/19	xe-0/0/19	Not supported on this port	Not supported on this port
20	Not supported on this port	ge-0/0/20	xe-0/0/20	Not supported on this port	Not supported on this port
21	Not supported on this port	ge-0/0/21	xe-0/0/21	Not supported on this port	Not supported on this port
22	Not supported on this port	ge-0/0/22	xe-0/0/22	Not supported on this port	Not supported on this port
23	Not supported on this port	ge-0/0/23	xe-0/0/23	Not supported on this port	Not supported on this port
24	Not supported on this port	ge-0/0/24	xe-0/0/24	Not supported on this port	Not supported on this port
25	Not supported on this port	ge-0/0/25	xe-0/0/25	Not supported on this port	Not supported on this port

Table 35: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
26	Not supported on this port	ge-0/0/26	xe-0/0/26	Not supported on this port	Not supported on this port
27	Not supported on this port	ge-0/0/27	xe-0/0/27	Not supported on this port	Not supported on this port
28	Not supported on this port	ge-0/0/28	xe-0/0/28	Not supported on this port	Not supported on this port
29	Not supported on this port	ge-0/0/29	xe-0/0/29	Not supported on this port	Not supported on this port
30	Not supported on this port	ge-0/0/30	xe-0/0/30	Not supported on this port	Not supported on this port
31	Not supported on this port	ge-0/0/31	xe-0/0/31	Not supported on this port	Not supported on this port
32	Not supported on this port	ge-0/0/32	xe-0/0/32	Not supported on this port	Not supported on this port
33	Not supported on this port	ge-0/0/33	xe-0/0/33	Not supported on this port	Not supported on this port
34	Not supported on this port	ge-0/0/34	xe-0/0/34	Not supported on this port	Not supported on this port
35	Not supported on this port	ge-0/0/35	xe-0/0/35	Not supported on this port	Not supported on this port
36	Not supported on this port	ge-0/0/36	xe-0/0/36	Not supported on this port	Not supported on this port
37	Not supported on this port	ge-0/0/37	xe-0/0/37	Not supported on this port	Not supported on this port
38	Not supported on this port	ge-0/0/38	xe-0/0/38	Not supported on this port	Not supported on this port

Table 35: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
39	Not supported on this port	ge-0/0/39	xe-0/0/39	Not supported on this port	Not supported on this port
40	Not supported on this port	ge-0/0/40	xe-0/0/40	Not supported on this port	Not supported on this port
41	Not supported on this port	ge-0/0/41	xe-0/0/41	Not supported on this port	Not supported on this port
42	fc-0/0/42	Not supported on this port	xe-0/0/42	Not supported on this port	Not supported on this port
43	fc-0/0/43	Not supported on this port	xe-0/0/43	Not supported on this port	Not supported on this port
44	fc-0/0/44	Not supported on this port	xe-0/0/44	Not supported on this port	Not supported on this port
45	fc-0/0/45	Not supported on this port	xe-0/0/45	Not supported on this port	Not supported on this port
46	fc-0/0/46	Not supported on this port	xe-0/0/46	Not supported on this port	Not supported on this port
47	fc-0/0/47	Not supported on this port	xe-0/0/47	Not supported on this port	Not supported on this port
Q0	Not supported on this port	Not supported on this port	xe-0/1/0 xe-0/1/1 xe-0/1/2 xe-0/1/3 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/0	xle-0/2/0

Table 35: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
Q1	Not supported on this port	Not supported on this port	xe-0/1/4 xe-0/1/5 xe-0/1/6 xe-0/1/7 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/1	xle-0/2/1
Q2	Not supported on this port	Not supported on this port	xe-0/1/8 xe-0/1/9 xe-0/1/10 xe-0/1/11 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/2	xle-0/2/2
Q3	Not supported on this port	Not supported on this port	xe-0/1/12 xe-0/1/13 xe-0/1/14 xe-0/1/15 NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/3	xle-0/2/3

Table 36: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
0	Not supported on this port	xe-0/0/0	Not supported on this port
1	Not supported on this port	xe-0/0/1	Not supported on this port
2	Not supported on this port	xe-0/0/2	Not supported on this port
3	Not supported on this port	xe-0/0/3	Not supported on this port
4	Not supported on this port	xe-0/0/4	Not supported on this port
5	Not supported on this port	xe-0/0/5	Not supported on this port
6	ge-0/0/6	xe-0/0/6	Not supported on this port
7	ge-0/0/7	xe-0/0/7	Not supported on this port
8	ge-0/0/8	xe-0/0/8	Not supported on this port
9	ge-0/0/9	xe-0/0/9	Not supported on this port
10	ge-0/0/10	xe-0/0/10	Not supported on this port
11	ge-0/0/11	xe-0/0/11	Not supported on this port
12	ge-0/0/12	xe-0/0/12	Not supported on this port
13	ge-0/0/13	xe-0/0/13	Not supported on this port
14	ge-0/0/14	xe-0/0/14	Not supported on this port
15	ge-0/0/15	xe-0/0/15	Not supported on this port
16	ge-0/0/16	xe-0/0/16	Not supported on this port
17	ge-0/0/17	xe-0/0/17	Not supported on this port
18	ge-0/0/18	xe-0/0/18	Not supported on this port

Table 36: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
19	ge-0/0/19	xe-0/0/19	Not supported on this port
20	ge-0/0/20	xe-0/0/20	Not supported on this port
21	ge-0/0/21	xe-0/0/21	Not supported on this port
22	ge-0/0/22	xe-0/0/22	Not supported on this port
23	ge-0/0/23	xe-0/0/23	Not supported on this port
24	ge-0/0/24	xe-0/0/24	Not supported on this port
25	ge-0/0/25	xe-0/0/25	Not supported on this port
26	ge-0/0/26	xe-0/0/26	Not supported on this port
27	ge-0/0/27	xe-0/0/27	Not supported on this port
28	ge-0/0/28	xe-0/0/28	Not supported on this port
29	ge-0/0/29	xe-0/0/29	Not supported on this port
30	ge-0/0/30	xe-0/0/30	Not supported on this port
31	ge-0/0/31	xe-0/0/31	Not supported on this port
32	ge-0/0/32	xe-0/0/32	Not supported on this port
33	ge-0/0/33	xe-0/0/33	Not supported on this port
34	ge-0/0/34	xe-0/0/34	Not supported on this port
35	ge-0/0/35	xe-0/0/35	Not supported on this port
36	ge-0/0/36	xe-0/0/36	Not supported on this port
37	ge-0/0/37	xe-0/0/37	Not supported on this port

Table 36: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
38	ge-0/0/38	xe-0/0/38	Not supported on this port
39	ge-0/0/39	xe-0/0/39	Not supported on this port
40	ge-0/0/40	xe-0/0/40	Not supported on this port
41	ge-0/0/41	xe-0/0/41	Not supported on this port
42	Not supported on this port	xe-0/0/42	Not supported on this port
43	Not supported on this port	xe-0/0/43	Not supported on this port
44	Not supported on this port	xe-0/0/44	Not supported on this port
45	Not supported on this port	xe-0/0/45	Not supported on this port
46	Not supported on this port	xe-0/0/46	Not supported on this port
47	Not supported on this port	xe-0/0/47	Not supported on this port
Q0	Not supported on this port	xe-0/1/0:0 xe-0/1/0:1 xe-0/1/0:2 xe-0/1/0:3	et-0/1/0
Q1	Not supported on this port	xe-0/1/1:0 xe-0/1/1:1 xe-0/1/1:2 xe-0/1/1:3	et-0/1/1

Table 36: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q2	Not supported on this port	xe-0/1/2:0 xe-0/1/2:1 xe-0/1/2:2 xe-0/1/2:3	et-0/1/2
Q3	Not supported on this port	xe-0/1/3:0 xe-0/1/3:1 xe-0/1/3:2 xe-0/1/3:3	et-0/1/3

Table 37: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	xe-0/0/0 xe-0/0/1 xe-0/0/2 xe-0/0/3	xle-0/1/0
Q1	xe-0/0/4 xe-0/0/5 xe-0/0/6 xe-0/0/7	xle-0/1/1
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	xle-0/1/2

Table 37: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q3	xe-0/0/12	xle-0/1/3
	xe-0/0/13	
	xe-0/0/14	
	xe-0/0/15	
Q4	xe-0/0/16	xle-0/1/4
	xe-0/0/17	
	xe-0/0/18	
	xe-0/0/19	
Q5	xe-0/0/20	xle-0/1/5
	xe-0/0/21	
	xe-0/0/22	
	xe-0/0/23	
Q6	xe-0/0/24	xle-0/1/6
	xe-0/0/25	
	xe-0/0/26	
	xe-0/0/27	
Q7	xe-0/0/28	xle-0/1/7
	xe-0/0/29	
	xe-0/0/30	
	xe-0/0/31	

Table 37: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q8	xe-0/0/32	xle-0/1/8
	xe-0/0/33	
	xe-0/0/34	
	xe-0/0/35	
Q9	xe-0/0/36	xle-0/1/9
	xe-0/0/37	
	xe-0/0/38	
	xe-0/0/39	
Q10	xe-0/0/40	xle-0/1/10
	xe-0/0/41	
	xe-0/0/42	
	xe-0/0/43	
Q11	xe-0/0/44	xle-0/1/11
	xe-0/0/45	
	xe-0/0/46	
	xe-0/0/47	
Q12	xe-0/0/48	xle-0/1/12
	xe-0/0/49	
	xe-0/0/50	
	xe-0/0/51	

Table 37: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q13	xe-0/0/52 xe-0/0/53 xe-0/0/54 xe-0/0/55	xle-0/1/13
Q14	xe-0/0/56 xe-0/0/57 xe-0/0/58 xe-0/0/59	xle-0/1/14
Q15	xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63	xle-0/1/15

Table 38: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q0	xe-0/0/0:0 xe-0/0/0:1 xe-0/0/0:2 xe-0/0/0:3	et-0/0/0
Q1	xe-0/0/1:0 xe-0/0/1:1 xe-0/0/1:2 xe-0/0/1:3	et-0/0/1

Table 38: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q2	xe-0/0/2:0 xe-0/0/2:1 xe-0/0/2:2 xe-0/0/2:3	et-0/0/2
Q3	xe-0/0/3:0 xe-0/0/3:1 xe-0/0/3:2 xe-0/0/3:3	et-0/0/3
Q4	xe-0/0/4:0 xe-0/0/4:1 xe-0/0/4:2 xe-0/0/4:3	et-0/0/4
Q5	xe-0/0/5:0 xe-0/0/5:1 xe-0/0/5:2 xe-0/0/5:3	et-0/0/5
Q6	xe-0/0/6:0 xe-0/0/6:1 xe-0/0/6:2 xe-0/0/6:3	et-0/0/6

Table 38: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q7	xe-0/0/7:0 xe-0/0/7:1 xe-0/0/7:2 xe-0/0/7:3	et-0/0/7
Q8	xe-0/0/8:0 xe-0/0/8:1 xe-0/0/8:2 xe-0/0/8:3	et-0/0/8
Q9	xe-0/0/9:0 xe-0/0/9:1 xe-0/0/9:2 xe-0/0/9:3	et-0/0/9
Q10	xe-0/0/10:0 xe-0/0/10:1 xe-0/0/10:2 xe-0/0/10:3	et-0/0/10
Q11	xe-0/0/11:0 xe-0/0/11:1 xe-0/0/11:2 xe-0/0/11:3	et-0/0/11

Table 38: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q12	xe-0/0/12:0 xe-0/0/12:1 xe-0/0/12:2 xe-0/0/12:3	et-0/0/12
Q13	xe-0/0/13:0 xe-0/0/13:1 xe-0/0/13:2 xe-0/0/13:3	et-0/0/13
Q14	xe-0/0/14:0 xe-0/0/14:1 xe-0/0/14:2 xe-0/0/14:3	et-0/0/14
Q15	xe-0/0/15:0 xe-0/0/15:1 xe-0/0/15:2 xe-0/0/15:3	et-0/0/15

Table 39: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	Not supported on this port	fte-0/1/0	xle-0/1/0
Q1	Not supported on this port	fte-0/1/1	xle-0/1/1

Table 39: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	fte-0/1/2	xle-0/1/2
Q3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15	fte-0/1/3	xle-0/1/3
Q4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19	fte-0/1/4	xle-0/1/4
Q5	xe-0/0/20 xe-0/0/21 xe-0/0/22 xe-0/0/23	fte-0/1/5	xle-0/1/5
Q6	xe-0/0/24 xe-0/0/25 xe-0/0/26 xe-0/0/27	fte-0/1/6	xle-0/1/6

Table 39: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q7	xe-0/0/28 xe-0/0/29 xe-0/0/30 xe-0/0/31	fte-0/1/7	xle-0/1/7
Q8	xe-0/0/32 xe-0/0/33 xe-0/0/34 xe-0/0/35	Not supported on this port	xle-0/1/8
Q9	xe-0/0/36 xe-0/0/37 xe-0/0/38 xe-0/0/39	Not supported on this port	xle-0/1/9
Q10	xe-0/0/40 xe-0/0/41 xe-0/0/42 xe-0/0/43	Not supported on this port	xle-0/1/10
Q11	xe-0/0/44 xe-0/0/45 xe-0/0/46 xe-0/0/47	Not supported on this port	xle-0/1/11

Table 39: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q12	xe-0/0/48 xe-0/0/49 xe-0/0/50 xe-0/0/51	Not supported on this port	xle-0/1/12
Q13	xe-0/0/52 xe-0/0/53 xe-0/0/54 xe-0/0/55	Not supported on this port	xle-0/1/13
Q14	xe-0/0/56 xe-0/0/57 xe-0/0/58 xe-0/0/59	Not supported on this port	xle-0/1/14
Q15	xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63	Not supported on this port	xle-0/1/15

Table 40: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
0	xe-0/0/0	Not supported on this port
1	xe-0/0/1	Not supported on this port

Table 40: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
2	xe-0/0/2	Not supported on this port
3	xe-0/0/3	Not supported on this port
4	xe-0/0/4	Not supported on this port
5	xe-0/0/5	Not supported on this port
6	xe-0/0/6	Not supported on this port
7	xe-0/0/7	Not supported on this port
8	xe-0/0/8	Not supported on this port
9	xe-0/0/9	Not supported on this port
10	xe-0/0/10	Not supported on this port
11	xe-0/0/11	Not supported on this port
12	xe-0/0/12	Not supported on this port
13	xe-0/0/13	Not supported on this port
14	xe-0/0/14	Not supported on this port
15	xe-0/0/15	Not supported on this port
16	xe-0/0/16	Not supported on this port
17	xe-0/0/17	Not supported on this port
18	xe-0/0/18	Not supported on this port
19	xe-0/0/19	Not supported on this port
20	xe-0/0/20	Not supported on this port

Table 40: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
21	xe-0/0/21	Not supported on this port
22	xe-0/0/22	Not supported on this port
23	xe-0/0/23	Not supported on this port
24	xe-0/0/24	Not supported on this port
25	xe-0/0/25	Not supported on this port
26	xe-0/0/26	Not supported on this port
27	xe-0/0/27	Not supported on this port
28	xe-0/0/28	Not supported on this port
29	xe-0/0/29	Not supported on this port
30	xe-0/0/30	Not supported on this port
31	xe-0/0/31	Not supported on this port
32	xe-0/0/32	Not supported on this port
33	xe-0/0/33	Not supported on this port
34	xe-0/0/34	Not supported on this port
35	xe-0/0/35	Not supported on this port
36	xe-0/0/36	Not supported on this port
37	xe-0/0/37	Not supported on this port
38	xe-0/0/38	Not supported on this port
39	xe-0/0/39	Not supported on this port

Table 40: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
40	xe-0/0/40	Not supported on this port
41	xe-0/0/41	Not supported on this port
42	xe-0/0/42	Not supported on this port
43	xe-0/0/43	Not supported on this port
44	xe-0/0/44	Not supported on this port
45	xe-0/0/45	Not supported on this port
46	xe-0/0/46	Not supported on this port
47	xe-0/0/47	Not supported on this port
48	xe-0/0/48:0 xe-0/0/48:1 xe-0/0/48:2 xe-0/0/48:3	et-0/0/48
49	xe-0/0/49:0 xe-0/0/49:1 xe-0/0/49:2 xe-0/0/49:3	et-0/0/49
50	xe-0/0/50:0 xe-0/0/50:1 xe-0/0/50:2 xe-0/0/50:3	et-0/0/50

Table 40: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
51	xe-0/0/51:0 xe-0/0/51:1 xe-0/0/51:2 xe-0/0/51:3	et-0/0/51
52	xe-0/0/52:0 xe-0/0/52:1 xe-0/0/52:2 xe-0/0/52:3	et-0/0/52
53	xe-0/0/53:0 xe-0/0/53:1 xe-0/0/53:2 xe-0/0/53:3	et-0/0/53

Table 41: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
0	xe-0/0/0	Not supported on this port	Not supported on this port
1	xe-0/0/1	Not supported on this port	Not supported on this port
2	xe-0/0/2	Not supported on this port	Not supported on this port
3	xe-0/0/3	Not supported on this port	Not supported on this port
4	xe-0/0/4	Not supported on this port	Not supported on this port

Table 41: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
5	xe-0/0/5	Not supported on this port	Not supported on this port
6	xe-0/0/6	Not supported on this port	Not supported on this port
7	xe-0/0/7	Not supported on this port	Not supported on this port
8	xe-0/0/8	Not supported on this port	Not supported on this port
9	xe-0/0/9	Not supported on this port	Not supported on this port
10	xe-0/0/10	Not supported on this port	Not supported on this port
11	xe-0/0/11	Not supported on this port	Not supported on this port
12	xe-0/0/12	Not supported on this port	Not supported on this port
13	xe-0/0/13	Not supported on this port	Not supported on this port
14	xe-0/0/14	Not supported on this port	Not supported on this port
15	xe-0/0/15	Not supported on this port	Not supported on this port
16	xe-0/0/16	Not supported on this port	Not supported on this port
17	xe-0/0/17	Not supported on this port	Not supported on this port
18	xe-0/0/18	Not supported on this port	Not supported on this port
19	xe-0/0/19	Not supported on this port	Not supported on this port
20	xe-0/0/20	Not supported on this port	Not supported on this port
21	xe-0/0/21	Not supported on this port	Not supported on this port
22	xe-0/0/22	Not supported on this port	Not supported on this port
23	xe-0/0/23	Not supported on this port	Not supported on this port

Table 41: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
24	xe-0/0/24	Not supported on this port	Not supported on this port
25	xe-0/0/25	Not supported on this port	Not supported on this port
26	xe-0/0/26	Not supported on this port	Not supported on this port
27	xe-0/0/27	Not supported on this port	Not supported on this port
28	xe-0/0/28	Not supported on this port	Not supported on this port
29	xe-0/0/29	Not supported on this port	Not supported on this port
30	xe-0/0/30	Not supported on this port	Not supported on this port
31	xe-0/0/31	Not supported on this port	Not supported on this port
32	xe-0/0/32	Not supported on this port	Not supported on this port
33	xe-0/0/33	Not supported on this port	Not supported on this port
34	xe-0/0/34	Not supported on this port	Not supported on this port
35	xe-0/0/35	Not supported on this port	Not supported on this port
36	xe-0/0/36	Not supported on this port	Not supported on this port
37	xe-0/0/37	Not supported on this port	Not supported on this port
38	xe-0/0/38	Not supported on this port	Not supported on this port
39	xe-0/0/39	Not supported on this port	Not supported on this port
40	xe-0/0/40	Not supported on this port	Not supported on this port
41	xe-0/0/41	Not supported on this port	Not supported on this port
42	xe-0/0/42	Not supported on this port	Not supported on this port

Table 41: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
43	xe-0/0/43	Not supported on this port	Not supported on this port
44	xe-0/0/44	Not supported on this port	Not supported on this port
45	xe-0/0/45	Not supported on this port	Not supported on this port
46	xe-0/0/46	Not supported on this port	Not supported on this port
47	xe-0/0/47	Not supported on this port	Not supported on this port
48	Not supported on this port	Not supported on this PIC	fte-0/1/0 NOTE: This interface is a fixed fte interface and cannot be changed to xle.
49	Not supported on this port	Not supported on this PIC	fte-0/1/1 NOTE: This interface is a fixed fte interface and cannot be changed to xle.
50	Not supported on this port	xle-0/1/2	fte-0/1/2 NOTE: By default, this interface is an fte interface but can be configured as an xle interface.
51	Not supported on this port	xle-0/1/3	fte-0/1/3 NOTE: By default, this interface is an fte interface but can be configured as an xle interface.

Table 41: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
52	Not supported on this port	xle-0/1/4 NOTE: By default, this interface is an xle interface but can be configured as an fte interface.	fte-0/1/4
53	Not supported on this port	xle-0/1/5 NOTE: By default, this interface is an xle interface but can be configured as an fte interface.	fte-0/1/5

Supported System Modes

NOTE: There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches depending on the system mode you configure. If you try to channelize ports that are restricted, the configuration is ignored.

The following system modes are available on the QFX5100-24Q switch:

- Default mode
- Mode-104-port
- Flexi-PIC mode
- Non-oversubscribed mode

See [Table 42 on page 196](#) for more information regarding the supported system modes for your switch.

The following system modes are available on the QFX5100-96S switch:

- Default-mode
- Non-oversubscribed mode

See [Table 42 on page 196](#) for more information regarding the supported system modes for your switch.

Table 42: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software

	Default-mode	Mode-104port	Flexi-pic-mode	Non-over-subscribed-mode
QFX5100-48S and QFX5100-48T	Not supported	Not supported	Not supported	Not supported
QFX5100-24Q	<p>Supported</p> <p>You do not need to configure the switch to be in this mode. On PIC 0, you can channelize all 24 40-Gbps QSFP+ ports. On PIC 1 and PIC 2, the 40-Gbps QSFP+ ports in the expansion modules are supported but cannot be channelized. In this mode, you can have one of two port combinations: 32 40-Gbps QSFP+ ports, or 96 10-Gigabit Ethernet ports plus 8 40-Gbps QSFP+ ports.</p>	<p>Supported</p> <p>On PIC 0, all 24 40-Gbps QSFP+ ports are channelized by default, which provides 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in an expansion module on PIC 1 are supported. On PIC 1, ports 0 and 2 are channelized by default, and ports 1 and 3 are disabled. If 40-Gbps QSFP+ ports contained in an expansion module are detected on PIC 2, they are ignored.</p>	<p>Supported</p> <p>On PIC 0, the first four ports (ports 0 through 3) cannot be channelized. 40-Gbps QSFP+ ports contained in expansion modules on PIC 1 and PIC 2 are supported but cannot be channelized.</p>	<p>Supported</p> <p>All 24 40-Gbps QSFP+ ports on PIC 0 can be channelized to 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in the expansion modules on PIC 1 and PIC 2 are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.</p>

Table 42: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software (*continued*)

	Default-mode	Mode-104port	Flexi-pic-mode	Non-over-subscribed-mode
QFX5100-96S	<p>Supported</p> <p>You do not need to configure the switch to be in this mode. On PIC 0, all 96 10-Gigabit Ethernet ports are supported. You can only channelize the 40-Gbps QSFP+ interfaces to 10-Gigabit Ethernet interfaces on ports 96 and 100. When you channelize the interfaces on ports 96 and 100, ports 97, 98, 99, 101, 102 and 103 are disabled.</p>	Not supported	Not supported	<p>Supported</p> <p>On PIC 0, all 96 10-Gigabit Ethernet ports are supported. However, the eight 40-Gbps QSFP+ ports are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.</p>

SEE ALSO

Rear Panel of a QFX3500 Device

Front Panel of a QFX3600 Device

Configuring the System Mode

You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q, QFX5100-96S, and QFX5110-32Q switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. By default, all QSFP+ interfaces are auto-channelized. Auto-channelization is not supported on interfaces contained in expansion modules or on Virtual Chassis ports. To disable auto-channelization, see [“Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches” on page 134](#) for more information.

NOTE: QFX5200 switches do not support System Mode.

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



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

See [Table 43 on page 198](#), [Table 44 on page 199](#), [Table 45 on page 200](#), and [Table 46 on page 201](#) for more information regarding the supported system modes for your switch.

Table 43: System Modes Supported on QFX5100 Switches with QFX-EM-4Q or QFX-PFA-4Q Expansion Modules Installed

	Default-mode	Mode-104port	Flexi-pic-mode	Non-over-subscribed mode
QFX5100-48S	Not supported	Not supported	Not supported	Not supported
QFX5100-24Q	Supported You do not need to configure the switch to be in this mode. On PIC 0, you can channelize all 24 40-Gbps QSFP+ ports. On PIC 1 and PIC 2, the 40-Gbps QSFP+ ports in the expansion modules are supported but cannot be channelized. In this mode, you can have one of two port combinations: 32 40-Gbps QSFP+ ports, or 96 10-Gigabit Ethernet ports plus 8 40-Gbps QSFP+ ports.	Supported On PIC 0, all 24 40-Gbps QSFP+ ports are channelized by default, which provides 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in an expansion module on PIC 1 are supported. On PIC 1, ports 0 and 2 are channelized by default, and ports 1 and 3 are disabled. If 40-Gbps QSFP+ ports contained in an expansion module are detected on PIC 2, they are ignored.	Supported On PIC 0, the first four ports (ports 0 through 3) cannot be channelized. 40-Gbps QSFP+ ports contained in expansion modules on PIC 1 and PIC 2 are supported but cannot be channelized.	Supported All 24 40-Gbps QSFP+ ports on PIC 0 can be channelized to 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in the expansion modules on PIC 1 and PIC 2 are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.

Table 43: System Modes Supported on QFX5100 Switches with QFX-EM-4Q or QFX-PFA-4Q Expansion Modules Installed (continued)

	Default-mode	Mode-104port	Flexi-pic-mode	Non-over-subscribed mode
QFX5100-96S	<p>Supported</p> <p>You do not need to configure the switch to be in this mode. On PIC 0, all 96 10-Gigabit Ethernet ports are supported. You can only channelize the 40-Gbps QSFP+ interfaces to 10-Gigabit Ethernet interfaces on ports 96 and 100. When you channelize the interfaces on ports 96 and 100, ports 97, 98, 99, 101, 102 and 103 are disabled.</p>	Not supported	Not supported	<p>Supported</p> <p>On PIC 0, all 96 10-Gigabit Ethernet ports are supported. However, the eight 40-Gbps QSFP+ ports are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.</p>

Table 44: System Modes Supported on QFX5100-24Q Switches with the EX4600-8F Expansion Module Installed

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

Table 45: System Modes Supported on QFX5100-24Q Switches with EX4600-8F and QFX-EM-4Q Expansion Modules Installed

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

Table 46: System Modes Supported on QFX5110-32Q Switches

	Default-mode	Mode-104port	Flexi-pic-mode	Non-over-subscribed-mode
QFX5110-32Q	<p>Supported</p> <p>All 32 QSFP+ ports on the switch are configured for 40-Gigabit Ethernet only. All ports are supported as access or uplink ports, but cannot be channelized.</p>	Not supported	<p>Supported</p> <ul style="list-style-type: none"> Ports 0 through 19 of the switch are configured for 40-Gigabit Ethernet and can be channelized to 4 independent 10-Gigabit Ethernet ports. Ports 20 through 27 are disabled. Ports 28 through 31 are configured as 100-Gigabit Ethernet. 	<p>Supported</p> <p>See QFX5110-32Q Port Panel for more details.</p>



CAUTION: Changing the system mode causes the FPC to reboot. Because there can be a slight loss of data while the FPC reboots, we recommend that you only configure the changes during a maintenance window for this release.



CAUTION: Take care when changing the channelization mode from Flexi-pic to default. If you have existing ports that are channelized in Flexi-pic mode, remove the channelization from the interface before changing the system mode. Changing the Flexi-pic mode to the default mode with channelized ports causes the ports to go down, log a system log error, and remain down. You must manually remove the channelization configuration on the ports to bring the ports up in default mode.

The following steps describe how to change the system mode.

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

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

For example:

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

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

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

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

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

Configuring the Port Type on QFX3600 Standalone Switches

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

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



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

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

Pipe write error: Broken pipe

flush operation failed

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

SEE ALSO

| [pic | 655](#)

Configuring the QSFP+ Port Type on QFX3500 Standalone Switches

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

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



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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

SEE ALSO

[Understanding Interface Naming Conventions](#) | 38

[pic](#) | 655

Configuring the QSFP+ Port Type on QFX5100 Devices

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

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

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



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

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

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

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

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

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

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

```
[edit chassis (QFX Series) node-group name node-device name pic 1]
```

```
user@switch# set fte port-range port-range-low port-range-high
```

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

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

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

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

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

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

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

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

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

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

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

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

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


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

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

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

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

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

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

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

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

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

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

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

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

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

SEE ALSO

Energy Efficient Ethernet Interfaces

IN THIS SECTION

- [Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces | 210](#)
- [Configuring Energy Efficient Ethernet on Interfaces | 211](#)
- [Verifying That EEE Is Saving Energy on Configured Ports | 212](#)

The energy efficient ethernet (EEE) helps in reducing the power consumption on physical layer devices. Configuring these EEE on interfaces includes enabling and EEE on Base-T copper ethernet port based on the power utilization and also verifying if EEE is saving energy on the configured ports.

Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces

Energy Efficient Ethernet (EEE), an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, reduces the power consumption of physical layer devices (PHYs) during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when the link is idle.

An Ethernet link consumes power even when a link is idle. EEE provides a method to utilize power in such a way that Ethernet links use power only during data transmission. EEE specifies a signaling protocol, Low Power Idle (LPI) for achieving the power saving during the idle time of Ethernet links. EEE allows PHYs to exchange LPI indications to signal the transition to low power mode when there is no traffic. LPI indicates when a link can go idle and when the link needs to resume after a predefined delay without impacting data transmission.

The following copper PHYs are standardized by IEEE 802.3az:

- 100BASE-T
- 1000BASE-T
- 10GBASE-T

Configuring Energy Efficient Ethernet on Interfaces

IN THIS SECTION

- [Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port | 211](#)
- [Disabling EEE on a Base-T Copper Ethernet Port | 211](#)

Energy Efficient Ethernet (EEE), an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, reduces the power consumption of physical layer devices (PHYs) during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when a link is idle.

NOTE: Configure EEE only on EEE-capable Base-T copper Ethernet ports. If you configure EEE on unsupported ports, the console displays the message: **“EEE not supported”**.

This topic describes:

Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port

To enable EEE on an EEE-capable Base-T copper Ethernet interface:

```
[edit]  
user@switch# set interfaces interface-name ether-options ieee-802-3az-eee
```

You can view the EEE status by using the **show interfaces *interface-name* detail** command.

Disabling EEE on a Base-T Copper Ethernet Port

To disable EEE on a Base-T copper Ethernet interface:

```
[edit]  
user@switch# delete interfaces interface-name ether-options ieee-802-3az-eee
```

By default, EEE is disabled on EEE-capable ports.

Verifying That EEE Is Saving Energy on Configured Ports

Purpose

Verify that enabling EEE saves energy on Base-T Copper Ethernet ports.

Action

You can see the amount of energy saved by EEE on an EX Series switch using the **show chassis power-budget-statistics** command.

1. View the power budget of an EX Series switch before enabling EEE.

- On an EX6210 switch:

```
user@switch>show chassis power-budget-statistics
```

```

      PSU  2      (EX6200-PWR-AC2500)      :    2500 W   Online
      PSU  3      )                        :         0 W   Offline
      Total Power supplied by all Online PSUs :    2500 W
      Power Redundancy Configuration          :    N+1
      Power Reserved for the Chassis           :    500 W

Fan Tray Statistics                        Base power  Power Used
FTC   0                        :    300 W         nan W
FPC Statistics                        Base power  Power Used   PoE
power  Priority
FPC   3   (EX6200-48T)           :    150 W         61.54 W
0 W      9
FPC   4   (EX6200-SRE64-4XS)     :    100 W         48.25 W
0 W      0
FPC   5   (EX6200-SRE64-4XS)     :    100 W         48.00 W
0 W      0
FPC   7   (EX6200-48T)           :    150 W         63.11 W
0 W      9
FPC   8   (EX6200-48T)           :    150 W         12.17 W
0 W      9

Total (non-PoE) Power allocated          :    950 W
Total Power allocated for PoE             :         0 W
Power Available (Redundant case)          :         0 W
Total Power Available                     :    1550 W

```

- On an EX4300 switch:

```
user@switch>show chassis power-budget-statistics fpc 1
```

```

PSU 1      (JPSU-1100-AC-AFO-A)          :    1100 W   Online
Power redundancy configuration           :    N+0
Total power supplied by all online PSUs :    1100 W
Base power reserved                      :    175 W
Non-PoE power being consumed             :    95 W
Total Power allocated for PoE            :    925 W
Total PoE power consumed                 :     0 W
Total PoE power remaining                :    925 W

```

2. Enable EEE on Base-T Copper Ethernet ports and save the configuration.

3. View the power budget of the switch after enabling EEE.

- On an EX6210 switch:

```
user@switch> show chassis power-budget-statistics
```

```

PSU 2      (EX6200-PWR-AC2500)          :    2500 W   Online
PSU 3      )                             :     0 W   Offline
Total Power supplied by all Online PSUs :    2500 W
Power Redundancy Configuration          :    N+1
Power Reserved for the Chassis           :    500 W

Fan Tray Statistics                      Base power  Power Used
FTC 0                                     :    300 W      nan W
FPC Statistics                          Base power  Power Used   PoE
power  Priority
FPC 3   (EX6200-48T)                    :    150 W      50.36 W
0 W      9
FPC 4   (EX6200-SRE64-4XS)              :    100 W      48.60 W
0 W      0
FPC 5   (EX6200-SRE64-4XS)              :    100 W      48.09 W
0 W      0
FPC 7   (EX6200-48T)                    :    150 W      51.38 W
0 W      9
FPC 8   (EX6200-48T)                    :    150 W      12.17 W
0 W      9

Total (non-PoE) Power allocated          :    950 W
Total Power allocated for PoE            :     0 W
Power Available (Redundant case)         :     0 W
Total Power Available                    :    1550 W

```

- On an EX4300 switch:

```
user@switch> show chassis power-budget-statistics fpc 1
```

```

PSU 1      (JPSU-1100-AC-AFO-A)      :    1100 W   Online
Power redundancy configuration        :    N+0
Total power supplied by all online PSUs :    1100 W
Base power reserved                   :    175 W
Non-PoE power being consumed          :    86 W
Total Power allocated for PoE         :    925 W
Total PoE power consumed              :    0 W
Total PoE power remaining             :    925 W

```

Meaning

On an EX6210 switch, the **Power Used** field in the output shows the actual power being consumed by the line card or SRE module, including PoE power. If you compare the values in the **Power Used** field before and after enabling EEE for FPC 3 and FPC 7, you will notice that power is saved when EEE is enabled.

NOTE: The **Power Used** field is displayed in the output only for EX6210 switches.

On an EX4300 switch, if you compare the values in the **Non-PoE power being consumed** field before and after enabling EEE, you will notice that power is saved when EEE is enabled.

Uplink Failure Detection

IN THIS SECTION

- [Overview of Uplink Failure Detection | 215](#)
- [Configuring Interfaces for Uplink Failure Detection | 217](#)
- [Example: Configuring Interfaces for Uplink Failure Detection | 218](#)
- [Verifying That Uplink Failure Detection Is Working Correctly | 224](#)

Uplink failure detection detects the failure on uplink interfaces and advertises this information to the downlink interfaces so that the switch over of interfaces is possible to avoid loss of traffic. The topics

below discuss the functions of uplink failure detections and the steps to configure and verify the working of it.

Overview of Uplink Failure Detection

IN THIS SECTION

- [Uplink Failure Detection Configuration | 215](#)
- [Failure Detection Pair | 216](#)
- [Debounce Interval | 217](#)

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate this information to the downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all of the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects disabled downlink interfaces, it switches over to the secondary link to help ensure that the traffic of the failed link is not dropped.


This topic describes:

Uplink Failure Detection Configuration

Uplink failure detection allows switches to monitor uplink interfaces to spot link failures. When a switch detects a link failure, it automatically disables the downlink interfaces bound to the uplink interface. A server that is connected to the disabled downlink interface triggers a network adapter failover to a secondary link to avoid any traffic loss.

[Figure 1 on page 216](#) illustrates a typical setup for uplink failure detection.

Figure 1: Uplink Failure Detection Configuration on Switches



For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down all downlink interfaces within that group. If any uplink interface returns to service, then the switch brings all downlink interfaces in that group back to service.

The switch can monitor both physical interface links and logical interface links for uplink failures, but you must put the two types of interfaces into separate groups.

NOTE: For logical interfaces, the server must send keepalives between the switch and the server to detect failure of logical links.

Failure Detection Pair

Uplink failure detection requires that you create pairs of uplink and downlink interfaces in a group. Each pair includes one of each of the following:

- A link-to-monitor interface—The link-to-monitor interfaces specify the uplinks the switch monitors. You can configure a maximum of 48 uplink interfaces as link-to-monitor interfaces for a group.
- A link-to-disable interface—The link-to-disable interfaces specify the downlinks the switch disables when the switch detects an uplink failure. You can configure a maximum of 48 downlinks to disable in the group.

The link-to-disable interfaces are bound to the link-to-monitor interfaces within the group. When a link-to-monitor interface returns to service, the switch automatically enables all link-to-disable interfaces in the group.

Debounce Interval

The debounce interval is the amount of time, in seconds, that elapses before the downlink interfaces are brought up after corresponding state changes of the uplink interfaces. You can configure the debounce interval for the uplink failure detection group. In absence of the debounce interval configuration, the downlink interfaces are brought up immediately after a state change of the uplink interfaces, which might introduce unnecessary state changes of the downlink interfaces, as well as unnecessary failovers on the servers connected to these ports.

In the event that the uplink interface goes down during the debounce interval, the debounce timer will start when the uplink interface comes back up. If the uplink interface goes down before the debounce interval expires, the debounce timer restarts when the uplink interface comes back up.

Any change you make to the debounce interval takes effect immediately. If you make a change to the debounce interval while the debounce timer is in effect, the change will take place if the new expiry time is in the future. If not, the timer stops immediately.

If uplink failure detection restarts during the debounce interval, the debounce timer resets, and the time that elapsed before uplink failure detection restarted is lost. The link-to-disable interface comes up without waiting for the debounce interval to elapse.

If the link-to-disable interface does not come up after the debounce timer expires, there might be latency between the time the debounce timer expires and the time when the link-to-disable interface actually comes up.

Configuring Interfaces for Uplink Failure Detection

You can configure uplink failure detection to help ensure balanced traffic flow. Using this feature, switches can monitor and detect link failure on uplink interfaces and can propagate the failure information to downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Follow these configuration guidelines:

- Configure an interface in only one group.
- Configure a maximum of 48 groups for each switch.
- Configure a maximum of 48 uplinks to monitor and a maximum of 48 downlinks to disable in each group.
- Configure physical links and logical links in separate groups.

To configure uplink failure detection on a switch:

1. Specify a name for an uplink failure detection group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name
```

2. Add an uplink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-monitor interface-name
```

3. Configure the debounce interval for the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group1 debounce-interval seconds
```

4. Repeat Step 2 for each uplink interface you add to the group.

5. Add a downlink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-disable interface-name
```

6. Repeat Step 4 for each downlink interface you add to the group.

NOTE: After you have configured an uplink failure detection group, use the [show uplink-failure-detection group \(Uplink Failure Detection\) group-name](#) command to verify that all interfaces in the group are up. If the interfaces are down, uplink failure detection does not work.

Example: Configuring Interfaces for Uplink Failure Detection

IN THIS SECTION

- [Requirements | 219](#)
- [Overview and Topology | 219](#)
- [Configuring Uplink Failure Detection on Both Switches | 220](#)
- [Verification | 222](#)

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate the failure information to the downlink interfaces. All of the network interface cards (NICs) on a server are configured as being either the primary link or the secondary link and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link to ensure that the traffic on the failed link is not dropped.

This example describes:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 19.2R1 or later for the QFX Series
- Two QFX5100, QFX5110, QFX5120, QFX5200, or QFX5210 switches
- Two aggregation switches
- One dual-homed server

Overview and Topology

The topology in this example illustrates how to configure uplink failure detection on Switch 1 and Switch B. Switch 1 and Switch 2 are both configured with a link-to-monitor interface (the uplink interface to the aggregation switch) and a link-to-disable interface (the downlink interface to the server). For simplicity, only one group of link-to-monitor interfaces and link-to-disable interfaces is configured for each switch. The server is dual-homed to both Switch 1 and Switch 2. In this scenario, if the link-to-monitor interface to Switch 1 is disabled, the server uses the link-to-monitor interface to Switch 2 instead.

NOTE: This example does not describe how to configure the dual-homed server or the aggregation switches. Please refer to the documentation for each of these devices for more information.

Figure 2 on page 220 illustrates a typical setup for uplink failure detection.

Figure 2: Uplink Failure Detection Configuration on Switches




Table 47 on page 220 lists uplink failure settings for each QFX3500 switch.

Table 47: Settings for Uplink Failure Protection Example

Switch 1	Switch 2
<ul style="list-style-type: none">• Group name: Group1• Link-to-monitor interface: xe-0/0/0• Link-to-disable interface: xe-0/0/1• Debounce interval: 20	<ul style="list-style-type: none">• Group name: Group2• Link-to-monitor interface: xe-0/0/0• Link-to-disable interface: xe-0/0/1• Debounce interval: 20

Configuring Uplink Failure Detection on Both Switches

To configure uplink failure detection on both switches, perform these tasks:

CLI Quick Configuration

To quickly configure uplink failure protection on Switch 1 and Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit protocols]

set uplink-failure-detection group group1

set uplink-failure-detection group group2

set uplink-failure-detection group group1 link-to-monitor xe-0/0/0
```

```

set uplink-failure-detection group group1 debounce-interval 20

set uplink-failure-detection group group2 link-to-monitor xe-0/0/0

set uplink-failure-detection group group2 debounce-interval 20

set uplink-failure-detection group group1 link-to-disable xe-0/0/1

set uplink-failure-detection group group2 link-to-disable xe-0/0/1

```

Step-by-Step Procedure

To configure uplink failure protection on both switches:

1. Specify a name for the uplink failure detection group on Switch 1:

```

[edit protocols]
user@switch# set uplink-failure-detection group group1

```

2. Add an uplink interface to the group on Switch 1:

```

[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-monitor xe-0/0/0

```

3. Add a downlink interface to the group on Switch 1:

```

[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-disable xe-0/0/1

```

4. Configure the debounce interval for group1 on Switch 1:

```

[edit protocols]
user@switch# set uplink-failure-detection group group1 debounce-interval 20

```

5. Specify a name for the uplink failure detection group on Switch 2:

```

[edit protocols]
user@switch# set uplink-failure-detection group group2

```

6. Add an uplink interface to the group on Switch 2:

```

[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-monitor xe-0/0/0

```

7. Configure the debounce interval for group2 on Switch 1:

```

[edit protocols]

```

```
user@switch# set uplink-failure-detection group group2 debounce-interval 20
```

8. Add a downlink interface to the group on Switch 2:

```
[edit protocols]
```

```
user@switch# set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

Results

Display the results of the configuration:

```
uplink-failure-detection {
  group {
    group1 {
      debounce-interval 20;
      link-to-monitor {
        xe-0/0/0;
      }
      link-to-disable {
        xe-0/0/1;
      }
    }
    group2 {
      debounce-interval 20;
      link-to-monitor {
        xe-0/0/0;
      }
      link-to-disable {
        xe-0/0/1;
      }
    }
  }
}
```

Verification

IN THIS SECTION

- [Verifying That Uplink Failure Detection is Working Correctly | 223](#)

To verify that uplink failure detection is working correctly, perform the following tasks on Switch 1 and Switch 2:

Verifying That Uplink Failure Detection is Working Correctly

Purpose

Verify that the switch disables the downlink interface when it detects an uplink failure.

Action

1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
```

```
Group           : group1
Uplink          : xe-0/0/0*
Downlink       : xe-0/0/1*
Failure Action  : Inactive
Debounce Interval : 20
```

NOTE: The asterisk (*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
```

```
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.
4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
```

```
Group           : group1
Uplink          : xe-0/0/0
Downlink       : xe-0/0/1
Failure Action  : Active
Debounce Interval : 20
```

Meaning

The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

Verifying That Uplink Failure Detection Is Working Correctly

Purpose

Verify that the switch disables the downlink interface when it detects an uplink failure.

Action

1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
```

```
Group           : group1
Uplink           : xe-0/0/0*
Downlink         : xe-0/0/1*
Failure Action   : Inactive
Debounce Interval : 20
```

NOTE: The asterisk (*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.
4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
```

```
Group           : group1
Uplink           : xe-0/0/0
Downlink         : xe-0/0/1
Failure Action   : Active
Debounce Interval : 20
```


Meaning

The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

Targeted Broadcast

IN THIS SECTION

- [Understanding Targeted Broadcast | 226](#)
- [Understanding IP Directed Broadcast | 227](#)
- [Configuring Targeted Broadcast | 229](#)
- [Example: Configuring IP Directed Broadcast on a Switch | 232](#)
- [Verifying IP Directed Broadcast Status | 237](#)

Targeted broadcast helps in remote administration tasks such as backups and wake-on LAN (WOL) on a LAN interface, and supports virtual routing and forwarding (VRF) instances. The below topic discuss the process and functioning of targeted broadcast, its configuration details, and the status of the broadcast on various platforms.

Understanding Targeted Broadcast

Targeted broadcast is a process of flooding a target subnet with Layer 3 broadcast IP packets originating from a different subnet. The intent of targeted broadcast is to flood the target subnet with the broadcast packets on a LAN interface without broadcasting to the entire network. Targeted broadcast is configured with various options on the egress interface of the router or switch and the IP packets are broadcast only on the LAN (egress) interface. Targeted broadcast helps you implement remote administration tasks such as backups and wake-on LAN (WOL) on a LAN interface, and supports virtual routing and forwarding (VRF) instances.

Regular Layer 3 broadcast IP packets originating from a subnet are broadcast within the same subnet. When these IP packets reach a different subnet, they are forwarded to the Routing Engine (to be forwarded to other applications). Because of this, remote administration tasks such as backups cannot be performed on a particular subnet through another subnet. As a workaround you can enable targeted broadcast, to forward broadcast packets that originate from a different subnet.

Layer 3 broadcast IP packets have a destination IP address that is a valid broadcast address for the target subnet. These IP packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. In the destination subnet, if the receiving router has targeted broadcast enabled on the egress interface, the IP packets are forwarded to an egress interface and the Routing Engine or to an egress interface only. The IP packets are then translated into broadcast IP packets which flood the target subnet only through the LAN interface (if there is no LAN interface, the packets are discarded), and all hosts on the target subnet receive the IP packets. If targeted broadcast is not enabled on the receiving router, the IP packets are treated as regular Layer 3 broadcast IP packets and are forwarded to the Routing Engine. If targeted broadcast is enabled without any options, the IP packets are forwarded to the Routing Engine.

Targeted broadcast can be configured to forward the IP packets only to an egress interface, which is helpful when the router is flooded with packets to process, or to both an egress interface and the Routing Engine.

NOTE: Targeted broadcast does not work when the targeted broadcast option **forward-and-send-to-re** and the traffic sampling option **sampling** are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of these options or enable the **sampling** option with the targeted broadcast option **forward-only** on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.

NOTE: Any firewall filter that is configured on the Routing Engine loopback interface (lo0) cannot be applied to IP packets that are forwarded to the Routing Engine as a result of a targeted broadcast. This is because broadcast packets are forwarded as flood next hop and not as local next hop traffic, and you can only apply a firewall filter to local next hop routes for traffic directed towards the Routing Engine.

SEE ALSO

| [targeted-broadcast](#) | [706](#)

Understanding IP Directed Broadcast

IN THIS SECTION

- [IP Directed Broadcast Overview](#) | [227](#)
- [IP Directed Broadcast Implementation](#) | [228](#)
- [When to Enable IP Directed Broadcast](#) | [228](#)
- [When Not to Enable IP Directed Broadcast](#) | [228](#)

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

This topic covers:

IP Directed Broadcast Overview

IP directed broadcast packets have a destination IP address that is a valid broadcast address for the subnet that is the target of the directed broadcast (the target subnet). The intent of an IP directed broadcast is to

flood the target subnet with the broadcast packets without broadcasting to the entire network. IP directed broadcast packets cannot originate from the target subnet.

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

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

IP Directed Broadcast Implementation

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

By default, IP directed broadcast is disabled.

When to Enable IP Directed Broadcast

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

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

When Not to Enable IP Directed Broadcast

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

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

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

SEE ALSO

[Example: Configuring IP Directed Broadcast on a Switch | 232](#)

Configuring IP Directed Broadcast for Switches

Example: Configuring IP Directed Broadcast on QFX Series Switch

Configuring Targeted Broadcast

IN THIS SECTION

- [Configuring Targeted Broadcast and Its Options | 229](#)
- [Display Targeted Broadcast Configuration Options | 231](#)

The following sections explain how to configure targeted broadcast on an egress interface and its options:

Configuring Targeted Broadcast and Its Options

You can configure targeted broadcast on an egress interface with different options. You can either allow the IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine or you can allow the IP packets to be forwarded on the egress interface only. Note that the packets are broadcast only if the egress interface is a LAN interface.

To configure targeted broadcast and its options:

1. Configure the physical interface.

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

2. Configure the logical unit number at the `[edit interfaces interface-name]` hierarchy level.

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

3. Configure the protocol family as inet at the `[edit interfaces interface-name unit interface-unit-number]` hierarchy level.

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

4. Configure targeted broadcast at the `[edit interfaces interface-name unit interface-unit-number family inet]` hierarchy level

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

5. Specify one of the following options as per requirement:

- To allow IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine.

```
[edit interfaces interface-name unit interface-unit-number family inet targeted-broadcast]
user@host# set forward-and-send-to-re
```

- To allow IP packets to be forwarded on the egress interface only.

```
[edit interfaces interface-name unit interface-unit-number family inet targeted-broadcast]
user@host# set forward-only
```

NOTE: Targeted broadcast does not work when the targeted broadcast option **forward-and-send-to-re** and the traffic sampling option **sampling** are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of these options or enable the **sampling** option with the targeted broadcast option **forward-only** on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.

Display Targeted Broadcast Configuration Options

IN THIS SECTION

- [Forward IP Packets On the Egress Interface and To the Routing Engine | 231](#)
- [Forward IP Packets On the Egress Interface Only | 231](#)

The following topics display targeted broadcast configuration with its various options:

Forward IP Packets On the Egress Interface and To the Routing Engine

Purpose

Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface and to send a copy of the IP packets to the Routing Engine.

Action

To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is ge-2/0/0, the unit value is set to 0, the protocol family is set to inet.

```
[edit interfaces interface-name unit interface-unit-number family inet]
user@host#show
targeted-broadcast {
    forward-and-send-to-re;
}
```

Forward IP Packets On the Egress Interface Only

Purpose

Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface only.

Action

To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is ge-2/0/0, the unit value is set to 0, the protocol family is set to inet.

```
[edit interfaces interface-name unit interface-unit-number family inet]
user@host#show
targeted-broadcast {
    forward-only;
}
```

SEE ALSO

| [targeted-broadcast](#) | [706](#)

Example: Configuring IP Directed Broadcast on a Switch

IN THIS SECTION

- [Requirements](#) | [232](#)
- [Overview and Topology](#) | [233](#)
- [Configuration](#) | [234](#)

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

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

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.4 or later for EX Series switches or Junos OS Release 15.1X53-D10 for QFX10000 switches.
- One PC
- One EX Series switch or QFX10000 switch

Before you configure IP directed broadcast for a subnet:

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

Overview and Topology

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

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

In this topology (see [Figure 3 on page 233](#)), a host is connected to an interface on a switch to manage the clients in subnet **10.1.2.1/24**. When the switch receives a packet with the broadcast IP address of the target subnet as its destination address, it forwards the packet to the subnet's Layer 3 interface and broadcasts it to all the hosts within the subnet.

Figure 3: Topology for IP Directed Broadcast




Table 48 on page 234 shows the settings of the components in this example.

Table 48: Components of the IP Directed Broadcast Topology

Property	Settings
Ingress VLAN name	v0
Ingress VLAN IP address	10.1.1.1/24
Egress VLAN name	v1
Egress VLAN IP address	10.1.2.1/24
Interfaces in VLAN v0	ge-0/0/3.0
Interfaces in VLAN v1	ge-0/0/0.0 and ge-0/0/1.0

Configuration

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

CLI Quick Configuration

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

```
[edit]
```

```
set interfaces ge-0/0/0.0 family ethernet-switching vlan members v1
```

```
set interfaces ge-0/0/1.0 family ethernet-switching vlan members v1
```

```
set interfaces irb.1 family inet address 10.1.2.1/24
```

```
set interfaces ge-0/0/3.0 family ethernet-switching vlan members v0
```

```
set interfaces vlan.0 family inet address 10.1.1.1/24
```

```
set vlans v1 l3-interface vlan.1
```

```
set vlans v0 l3-interface irb.0
```

```
set interfaces vlan.1 family inet targeted-broadcast
```

Step-by-Step Procedure

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

1. Add logical interface **ge-0/0/0.0** to VLAN **v1**:

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

2. Add logical interface **ge-0/0/1.0** to VLAN **v1**:

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

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

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

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

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

5. Configure the IP address for the ingress VLAN:

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

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

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

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

```
[edit interfaces]
user@switch# set vlan.1 family inet targeted-broadcast
user@switch# set vlan.0 family inet targeted-broadcast
```

Results

Check the results:

```
user@switch# show
interfaces {
```

```

ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members v1;
      }
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members v1;
      }
    }
  }
}
ge-0/0/3 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members v0;
      }
    }
  }
}
vlan {
  unit 0 {
    family inet {
      targeted-broadcast;
      address 10.1.1.1/24;
    }
  }
  unit 1 {
    family inet {
      targeted-broadcast;
      address 10.1.2.1/24;
    }
  }
}
vlands {
  default;
  v0 {

```

```

    l3-interface vlan.0;
  }
  v1 {
    l3-interface vlan.1;
  }
}

```

SEE ALSO

| *Configuring IP Directed Broadcast for Switches*

Verifying IP Directed Broadcast Status

Purpose

Verify that IP directed broadcast is enabled and is working on the subnet.

Action

Use the **show vlans extensive** command to verify that IP directed broadcast is enabled and working on the subnet as shown in [“Example: Configuring IP Directed Broadcast on a Switch” on page 232.](#)

ARP

IN THIS SECTION

- [Static ARP Table Entries Overview | 238](#)
- [Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses | 238](#)
- [Restricted and Unrestricted Proxy ARP Overview | 240](#)
- [Configuring Restricted and Unrestricted Proxy ARP | 242](#)
- [Configuring Gratuitous ARP | 243](#)

Static address resolution protocol (ARP) table entries are repended to by default when the destination address of the ARP is on the local network. These static ARP addresses can be configured for Ethernet or

Gigabit Ethernet interfaces. The topics below discuss the overview of static ARP table entries, restricted and unrestricted proxy ARP, configuration details to map the IP addresses to the MAC addresses.

Static ARP Table Entries Overview

For Fast Ethernet, Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure static ARP table entries, defining mappings between IP and MAC addresses.

SEE ALSO

Ethernet Interfaces User Guide for Routing Devices

Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses

By default, the device responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the device to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

Also, unlike dynamically learned ARP entries, static ARP entries do not age out. You can also configure static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces interface-name unit logical-unit-number family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

To configure static ARP entries:

1. In the configuration mode, at the **[edit]** hierarchy level, configure the router interface on which the ARP table entries for the router is configured.

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

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

NOTE: When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

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

3. Configure a static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. 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
```

4. Configure another static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. You can also associate a multicast MAC address with a unicast IP address by including the **multicast-mac** option with the **arp** statement. You can optionally configure the router to respond to ARP requests for the specified IP address by using the **publish** option with the **arp** statement.

NOTE: For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

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

NOTE: The Junos OS supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

SEE ALSO

[arp | 715](#)

Management Ethernet Interface Overview

Applying Policers

Configuring an Unnumbered Interface

Ethernet Interfaces User Guide for Routing Devices

Restricted and Unrestricted Proxy ARP Overview

By default, the Junos OS responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is local to the incoming interface.

For Ethernet Interfaces, you can configure the router or switches to proxy-reply to the ARP requests using the restricted or unrestricted proxy ARP configuration.

You might want to configure restricted or unrestricted proxy ARP for routers that act as provider edge (PE) devices in Ethernet Layer 2 LAN switching domains.

NOTE: From Junos OS Release 10.0 onward, Junos OS does not respond to proxy ARP requests with the default route 0.0.0.0. This behavior is in compliance with RFC 1027.

Restricted Proxy ARP

Restricted proxy ARP enables the router or switch to respond to the ARP requests in which the physical networks of the source and target are not the same and the router or switch has an active route to the target address in the ARP request. The router does not reply if the target address is on the same subnet and the same interface as the ARP requestor.

Unrestricted Proxy ARP

Unrestricted proxy ARP enables the router or switch to respond to any ARP request, on condition that the router has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route.



WARNING: If you configure unrestricted proxy ARP, the proxy router replies to ARP requests for the target IP address on the same interface as the incoming ARP request. This behavior is appropriate for cable modem termination system (CMTS) environments, but might cause Layer 2 reachability problems if you enable unrestricted proxy ARP in other environments.

When an IP client broadcasts the ARP request across the Ethernet wire, the end node with the correct IP address responds to the ARP request and provides the correct MAC address. If the unrestricted proxy ARP feature is enabled, the router response is redundant and might fool the IP client into determining that the destination MAC address within its own subnet is the same as the address of the router.

NOTE: While the destination address can be remote, the source address of the ARP request must be on the same subnet as the interface upon which the ARP request is received. For security reasons, this rule applies to both unrestricted and restricted proxy ARP.

Topology Considerations for Unrestricted Proxy ARP

In most situations, you should not configure the router or switch to perform unrestricted proxy ARP. Do so only for special situations, such as when cable modems are used. [Figure 4 on page 242](#) and [Figure 5 on page 242](#) show examples of situations in which you might want to configure unrestricted proxy ARP.

In [Figure 4 on page 242](#), the edge device is not running any IP protocols. In this case, you configure the core router to perform unrestricted proxy ARP. The edge device is the client of the proxy.

In [Figure 5 on page 242](#), the Broadband Remote Access Server (B-RAS) routers are not running any IP protocols. In this case, you configure unrestricted proxy ARP on the B-RAS interfaces. This allows the core device to behave as though it is directly connected to the end users.

Figure 4: Edge Device Case for Unrestricted Proxy ARP





Figure 5: Core Device Case for Unrestricted Proxy ARP



SEE ALSO

| *Ethernet Interfaces User Guide for Routing Devices*

Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

```
proxy-arp (restricted |unrestricted);
```

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*]

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.

NOTE: When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the **no-gratuitous-arp-reply** statement. See [“Configuring Gratuitous ARP” on page 243](#) for information about how to disable responses to gratuitous ARP requests.

SEE ALSO

| *Ethernet Interfaces User Guide for Routing Devices*

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests help detect duplicate IP addresses. A gratuitous ARP is a broadcast request for a router’s own IP address. If a router or switch sends an ARP request for its own IP address and no ARP replies are received, the router- or switch-assigned IP address is not being used by other nodes. However, if a router or switch sends an ARP request for its own IP address and an ARP reply is received, the router- or switch-assigned IP address is already being used by another node.

Gratuitous ARP replies are reply packets sent to the broadcast MAC address with the target IP address set to be the same as the sender's IP address. When the router or switch receives a gratuitous ARP reply, the router or switch can insert an entry for that reply in the ARP cache. By default, updating the ARP cache on gratuitous ARP replies is disabled on the router or switch.

To enable updating of the ARP cache for gratuitous ARPs:

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

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

2. Include the **gratuitous-arp-reply** statement.

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

To restore the default behavior, that is, to disable updating of the ARP cache for gratuitous ARP, delete the **gratuitous-arp-reply** statement from the configuration:

```
[edit interfaces interface-name]
user@host# delete gratuitous-arp-reply;
```

By default, the router or switch responds to gratuitous ARP requests. However, on Ethernet interfaces, you can disable responses to gratuitous ARP requests.

To disable responses to gratuitous ARP requests:

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

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

2. Include the **no-gratuitous-arp-request** statement.

```
[edit interfaces interface-name]
user@host# set no-gratuitous-arp-request
```

To return to the default—that is, to respond to gratuitous ARP requests—delete the **no-gratuitous-arp-request** statement from the configuration:

```
[edit interfaces interface-name]
user@host# delete no-gratuitous-arp-request
```

SEE ALSO

[*gratuitous-arp-reply*](#)

[*no-gratuitous-arp-request*](#)

[*Ethernet Interfaces Overview*](#)

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

Resilient Hashing on LAGs and ECMP groups

IN THIS SECTION

- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in LAGs/ECMP Groups | 245](#)
- [Configuring Resilient Hashing for LAGs/ECMP Groups | 248](#)

Resilient hashing helps minimize the flow remapping across equal cost multipath (ECMP) groups and LAGs in a load-balanced system. The topics below discuss the working, usage and configuring of resilient hashing on link aggregation groups (LAGs) and ECMP groups.

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

IN THIS SECTION

- [Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing | 246](#)
- [Limitations and Caveats for Resilient Hashing | 247](#)

- [Resilient Hashing on LAGs | 247](#)
- [Resilient Hashing on ECMP | 248](#)

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

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

Resilient hashing works in conjunction with the default static hashing algorithm. When members are added to or deleted from a LAG/ECMP group, the static hashing algorithm might remap destination paths. With resilient hashing, the chances of a flow being remapped are minimal if its path is unaffected by the LAG/ECMP group's member change. When a flow is affected by a member change, the Packet Forwarding Engine rebalances the flow by reprogramming the flow set table.

Resilient hashing thus provides the following benefits:

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

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

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

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

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

Table 49: Destination Path Results for Static Hashing and for Resilient Hashing When Members Are Added to or Deleted from LAGs

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

Limitations and Caveats for Resilient Hashing

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

- Resilient hashing applies only to unicast traffic.
- Resilient hashing supports a maximum of 1024 LAGs, with each group having a maximum of 256 members.
- Resilient hashing does not guarantee that traffic distribution is even across all group members—it depends on the traffic pattern and on the organization of the resilient hashing flow set table in hardware. Resilient hashing *minimizes* remapping of flows to destination links when members are added to or deleted from the group.
- If resilient hashing is enabled on a LAG or ECMP group and if **set forwarding-options enhanced-hash-key** with one of the options **hash-mode**, **inet**, **inet6**, or **layer2** is used, some flows might change destination links, because the new hash parameters might generate new hash indexes for the flows, and hence the new destination links.
- Resilient hashing is not supported on Virtual Chassis port (VCP) links.
- LAG-based resilient hashing is not supported on QFX5200 and QFX5210 switches. ECMP-based resilient hashing is supported on those switches.

Resilient Hashing on LAGs

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

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

Resilient Hashing on ECMP

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

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

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

Configuring Resilient Hashing for LAGs/ECMP Groups

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

This topic includes:

1. [Configuring Resilient Hashing on LAGs | 248](#)
2. [Configuring Resilient Hashing on ECMP Groups | 249](#)

Configuring Resilient Hashing on LAGs

NOTE: LAG-based resilient hashing is not supported on QFX5200 and QFX5210 switches. ECMP-based resilient hashing is supported on those switches.

To enable resilient hashing for a LAG:

- Configure resilient hashing on the aggregated Ethernet interface:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options resilient-hash
```

- (Optional) Configure a specific value for the resilient-hash seed. This value will apply only to the HASH2 engine:

```
[edit]
```



```
user@switch# set forwarding-options enhanced-hash-key resilient-hash-seed seed-value
```

Configuring Resilient Hashing on ECMP Groups

To enable resilient hashing for ECMP groups:

- Configure resilient hashing for ECMP:

```
[edit forwarding-options]  
user@switch# set enhanced-hash-key ecmp-resilient-hash
```

NOTE: When resilient hashing is added or removed, the traffic distribution across all members of an ECMP group for a given flow are reprogrammed and, as a result, some flows might be remapped to new ECMP group members.

Generic Routing Encapsulation (GRE)

IN THIS SECTION

- [Understanding Generic Routing Encapsulation | 250](#)
- [Configuring Generic Routing Encapsulation Tunneling | 255](#)
- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly | 256](#)

Generic routing encapsulation (GRE) provides a secure path for transporting packets of data by tunneling the packets. The below topics discuss the tunneling of GRE, encapsulation and de-capsulation process, configuring GREs and verifying the working of GREs.

Understanding Generic Routing Encapsulation

IN THIS SECTION

- Overview of GRE | 250
- GRE Tunneling | 251
- Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100, QFX10000, and OCX Series Switches | 253
- Configuration Limitations | 254

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

This topic describes:

Overview of GRE

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

GRE is described in RFC 2784 (obsoletes earlier RFCs 1701 and 1702). The switches support RFC 2784, but not completely. (For a list of limitations, see [“Configuration Limitations” on page 254.](#))

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

GRE Tunneling

IN THIS SECTION

- [Encapsulation and De-Encapsulation on the Switch | 251](#)
- [Number of Source and Destination Tunnels Allowed on a Switch | 252](#)
- [Class of Service on GRE Tunnels | 252](#)
- [Applying Firewall Filters to GRE Traffic | 252](#)

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

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

For details about GRE tunneling, see:

Encapsulation and De-Encapsulation on the Switch

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

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

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

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

Number of Source and Destination Tunnels Allowed on a Switch

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

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

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

Class of Service on GRE Tunnels

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

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

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

Applying Firewall Filters to GRE Traffic

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a switch. (For details, see [Firewall Filters for EX Series Switches Overview](#).) Because of the encapsulation and de-encapsulation performed by GRE, you are constrained as to where you can apply a firewall filter to filter tunneled packets and which header will be affected. [Table 50 on page 253](#) identifies these constraints.

Table 50: Firewall Filter Application Points for Tunneled Packets

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

Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100, QFX10000, and OCX Series Switches

You can also use a firewall filter to de-encapsulate GRE traffic on switches . This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. See *Configuring a Firewall Filter to De-Encapsulate GRE Traffic* for information about how to configure a firewall filter for this purpose.

Configuration Limitations

Table 51 on page 254 lists features that are not supported with GRE.

Table 51: Features Not Supported with GRE

EX Switches	QFX Switches
MPLS over GRE tunnels	MPLS over GRE tunnels
GRE keepalives	GRE keepalives
GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets	GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets
BGP dynamic tunnels	BGP dynamic tunnels
Outer IP address must be IPv4	Outer IP address must be IPv4
Virtual routing instances	On QFX10002 , QFX10008 and QFX5K Series switches, If you configure GRE tunneling with the underlying ECMP next-hop instead of a Unicast next-hop, GRE tunnel encapsulation fails and network traffic is dropped
Bidirectional Forwarding Detection (BFD) protocol over GRE distributed mode	
OSPF limitation—Enabling OSPF on a GRE interface creates two equal-cost routes to the destination: one through the Ethernet network or uplink interface and the other through the tunnel interface. If data is routed through the tunnel interface, the tunnel might fail. To keep the interface operational, we recommend that you use a static route, disable OSPF on the tunnel interface, or configure the peer not to advertise the tunnel destination over the tunnel interface.	
	QFX series switches do not support configuring GRE interface and the underlying tunnel source interface in two different routing instances. If you try this configuration, it will result in a commit error.

SEE ALSO

Configuring Generic Routing Encapsulation Tunneling

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

You can also use a firewall filter to de-encapsulate GRE traffic on QFX5100 and OCX Series switches. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. For more information on this feature, see *Configuring a Firewall Filter to De-Encapsulate GRE Traffic*.

To configure a GRE tunnel port on a switch:

1. Determine the network port or uplink port on your switch to convert to a GRE tunnel port.
2. Configure the port as a tunnel port for GRE tunnel services:

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

This topic describes:

1. [Configuring a GRE Tunnel | 255](#)

Configuring a GRE Tunnel

To configure a GRE tunnel interface:

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

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

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

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

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

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

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

3. Specify the destination address:

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

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

NOTE: On QFX10002 and QFX10008 switches, If you configure GRE tunneling with the underlying ECMP next-hop instead of Unicast next-hop, GRE tunnel encapsulation fails and the network traffic is dropped.

NOTE: Indirect egress next-hops is currently not supported in the GRE implementation for QFX10000 switches.

SEE ALSO

| [Configuring a Firewall Filter to De-Encapsulate GRE Traffic](#)

Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

Purpose

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

Action

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

```
user@switch> show interfaces gr-0/0/0.0
```

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

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
    Flags: None
    Addresses, Flags: Is-Primary
      Local: 10.0.0.0
```

Meaning

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

Understanding Per-Packet Load Balancing


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 re-chosen using the hash algorithm. Starting in Junos OS Release 18.3R1, for MX series routers, the default behavior for IPv6, GRE, and PPPoE packet hash computation was modified to include the flow-label field for improved load-balancing in certain cases (you can use the **no-payload** option to revert to the previous method for hash computation). See *Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers* for details.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This feature is called *per-packet load balancing*. The naming may be counter-intuitive. However, Junos *per-packet* load balancing is functionally equivalent to what other vendors may term *per-flow* load balancing. You can use load balancing to spread traffic across multiple paths between routers.

Figure 6 on page 258 shows a simple load balancing scenario. Device R1 is in AS 64500 and is connected to both Device R2 and Device R3, which are in AS 64501. Device R1 can be configured to load balance traffic across the two links.

Figure 6: Simple Load Balancing Scenario



Starting in Junos OS 13.3R3, for MX Series 5G Universal Routing Platforms with modular port concentrators (MPCs) only, you can configure consistent load balancing, which prevents the reordering of all flows to active paths in an equal-cost multipath (ECMP) group when one or more next-hop paths fail. Only flows for paths that are inactive are redirected to another active next-hop path. Flows mapped to servers that remain active are maintained. This feature applies only to external BGP peers.

Starting in Junos OS Release 19.1R1, on QFX10000 switches, you can configure load balancing of IPv4 or IPv6 packets by using GPRS Tunneling Protocol-tunnel endpoint identifier (GTP-TEID) field hash calculations. The GTP-TEID hashing is added to the Layer 2 and Layer 3 field hashing that you have already configured. To enable this feature on QFX10000 switches, configure the **gtp-tunnel-endpoint-identifier** statement at the **[edit forwarding-options enhanced-hash-key family inet]** or the **[edit forwarding-options enhanced-hash-key family inet6]** hierarchy Level. GTP versions 1 and 2 are supported; they support only user data. You must use UDP port number 2152 for both GTP versions.

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1, on QFX10000 switches, you can configure load balancing of IPv4 or IPv6 packets by using GPRS Tunneling Protocol-tunnel endpoint identifier (GTP-TEID) field hash calculations
18.3R1	Starting in Junos OS Release 18.3R1, for MX series routers, the default behavior for IPv6, GRE, and PPPoE packet hash computation was modified to include the flow-label field for improved load-balancing in certain cases (you can use the no-payload option to revert to the previous method for hash computation).
13.3R3	Starting in Junos OS 13.3R3, for MX Series 5G Universal Routing Platforms with modular port concentrators (MPCs) only, you can configure consistent load balancing, which prevents the reordering of all flows to active paths in an equal-cost multipath (ECMP) group when one or more next-hop paths fail.

RELATED DOCUMENTATION

Example: Load Balancing BGP Traffic

Configuring Per-Packet Load Balancing

Configuring Load Balancing Based on MPLS Labels

Configuring Load Balancing for Ethernet Pseudowires

Configuring Load Balancing Based on MAC Addresses

Configuring VPLS Load Balancing Based on IP and MPLS Information

Configuring VPLS Load Balancing on MX Series 5G Universal Routing Platforms

Configuring Consistent Load Balancing for ECMP Groups

2

CHAPTER

Configuring Aggregated Ethernet Interfaces

[Aggregated Ethernet Interfaces | 261](#)

[Load Balancing for Aggregated Ethernet Interfaces | 352](#)

Aggregated Ethernet Interfaces

IN THIS SECTION

- [Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)
- [Forcing LAG Links or Interfaces with Limited LACP Capability to Be Up | 267](#)
- [Configuring an Aggregated Ethernet Interface | 267](#)
- [Configuring Tagged Aggregated Ethernet Interfaces | 268](#)
- [Configuring Untagged Aggregated Ethernet Interfaces | 270](#)
- [Configuring the Number of Aggregated Ethernet Interfaces on the Device \(Enhanced Layer 2 Software\) | 271](#)
- [Example: Configuring Aggregated Ethernet Interfaces | 271](#)
- [Deleting an Aggregated Ethernet Interface | 273](#)
- [Understanding Local Link Bias | 274](#)
- [Configuring Local Link Bias | 276](#)
- [Understanding Local Minimum Links | 276](#)
- [Troubleshooting an Aggregated Ethernet Interface | 279](#)
- [Configuring Link Aggregation | 281](#)
- [Configuring Aggregated Ethernet Link Protection | 286](#)
- [Configuring Aggregated Ethernet Link Speed | 288](#)
- [Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface | 290](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch | 299](#)
- [Configuring Aggregated Ethernet LACP | 305](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)
- [Configuring LACP Hold-UP Timer to Prevent Link Flapping on LAG Interfaces | 319](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets | 320](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 321](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch | 328](#)
- [Understanding Independent Micro BFD Sessions for LAG | 333](#)
- [Configuring Micro BFD Sessions for LAG | 336](#)

- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic | 342](#)
- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) | 349](#)

The below topics discuss the overview aggregated ethernet interfaces, configuration details of link aggregation and aggregated Ethernet interfaces, troubleshooting and verification of aggregated Ethernet Interfaces.

Understanding Aggregated Ethernet Interfaces and LACP for Switches

IN THIS SECTION

- [Link Aggregation Group | 263](#)
- [Link Aggregation Control Protocol \(LACP\) | 265](#)

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*.

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. 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: On QFX5100, EX4600, QFX10002 standalone switches, and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.

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

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

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

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

Link Aggregation Group

You configure a LAG by specifying the link number as a physical device and then associating a set of interfaces (ports) with the link. All the interfaces must have the same speed and be in full-duplex mode. Juniper Networks Junos operating system (Junos OS) for EX Series Ethernet Switches assigns a unique ID and port priority to each interface. The ID and priority are not configurable.

The number of interfaces that can be grouped into a LAG and the total number of LAGs supported on a switch varies according to switch model. [Table 52 on page 263](#) lists the EX Series switches and the maximum number of interfaces per LAG and the maximum number of LAGs they support.

NOTE: For Junos OS Evolved, the software does not impose a limit 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.

Table 52: Maximum Interfaces per LAG and Maximum LAGs per Switch

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX2200	8	32
EX2300	8	128
EX3200	8	32
EX3300 and EX3300 Virtual Chassis	8	32

Table 52: Maximum Interfaces per LAG and Maximum LAGs per Switch (*continued*)

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX3400	16	128
EX4200 and EX4200 Virtual Chassis	8	111
EX4300 and EX4300 Virtual Chassis	16	128
EX4500, EX4500 Virtual Chassis, EX4550, and EX4550 Virtual Chassis	8	111
EX4600	32	128
EX6200	8	111
EX8200	12	255
EX8200 Virtual Chassis	12	239
EX9200	64	150

To create a LAG:

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

Keep in mind these hardware and software guidelines:

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

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

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and up to 448 LAGs are supported on QFX3500, QFX3600, EX4600, and OCX Series switches, and up to 1,000 LAGs are supported on QFX5100, QFX5200, QFX5110, QFX10002, QFX10008, and QFX10016 switches.

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

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and In a Junos Fusion, up to 1,000 LAGs are supported on QFX10002 switches acting as aggregation devices.
- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed and be in full-duplex mode.

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

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

Link Aggregation Control Protocol (LACP)

LACP is one method of bundling several physical interfaces to form one logical aggregated Ethernet interface. By default, Ethernet links do not exchange LACP protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit LACP PDUs,

or you can configure the links to passively transmit them, sending out LACP PDUs only when the Ethernet link receives them from the remote end. The LACP mode can be active or passive. The transmitting link is known as the *actor*, and the receiving link is known as the *partner*. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode. You can configure both VLAN-tagged and untagged aggregated Ethernet interfaces without LACP enabled. LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

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

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

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

When LACP is configured, it detects misconfigurations on the local end or the remote end of the link. Thus, LACP can help prevent communication failure:

- When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail.
- When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

SEE ALSO

| *Verifying the Status of a LAG Interface*

Forcing LAG Links or Interfaces with Limited LACP Capability to Be Up

A link without Link Access Control Protocol (LACP) configuration remains down and cannot be accessed by the provider edge (PE) devices in the topology. You can configure the force-up feature in LACP on a PE device for which you need connectivity.

To ensure that the peer with limited LACP capability is up and accessible on the LAG network, configure one of the aggregated Ethernet links or interfaces on a PE device to be up by using the appropriate hierarchy level on your device:

- **set interfaces *interface-name* ether-options 802.3ad lacp force-up**
- **set interfaces *interface-name* aggregated-ether-options lacp force-up**

In a standalone or a virtual chassis environment configured with Aggregated Ethernet (AE) :

- if an aggregated Ethernet interface (AE) on a switch has multiple member links and one member link in that AE is in the force-up state with its peer's LACP down, and then if LACP comes up partially—that is, if LACP is established with a non-force-up member link—force-up is disabled on the member link on which force-up has been set, and that member link is ready for connection establishment through LACP. Force-up is eligible only if the server-side interface has LACP issues.

Configuring an Aggregated Ethernet Interface

You can associate a physical interface with an aggregated Ethernet interface.

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 ether-options 802.3ad aex
```

You specify the interface instance number *x* to complete the link association; *x* can be from 0 through 480, for a total of 480 aggregated interfaces on MX Series routers or EX9200 switches. 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*.

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.

SEE ALSO

Aggregated Ethernet Interfaces Overview

Ethernet Interfaces User Guide for Routing Devices

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](#) | **681**

Configuring Untagged Aggregated Ethernet Interfaces

When you configure an untagged Aggregated Ethernet interface, the existing rules for untagged interfaces apply. These rules are as follows:

- 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.
- You cannot include the **vlan-id** statement in the configuration of the logical interface.

Configure an untagged aggregated Ethernet interface by omitting the **vlan-tagging** and **vlan-id** statements from the configuration:

```
[edit interfaces]
ge-1/1/1 {
  ether-options {
    802.3ad ae0;
  }
}
ae0 {
  # vlan-tagging; OMIT FOR UNTAGGED AE CONFIGURATIONS
  unit 0 {
    # vlan-id 100; OMIT FOR UNTAGGED AE CONFIGURATIONS
    family inet {
      address 10.0.0.1/24 {
        vrrp-group 0 {
          virtual-address 192.168.110.0;
          priority 200;
        }
      }
    }
  }
}
```

SEE ALSO

Ethernet Interfaces User Guide for Routing Devices

Configuring the Number of Aggregated Ethernet Interfaces on the Device (Enhanced Layer 2 Software)

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.

On MX Series routers and EX9200 switches, you can configure a maximum of 480 aggregated interfaces. The aggregated interfaces (LAG bundles) are numbered from **ae0** through **ae479** on MX Series routers and EX9200 switches.

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* ether-options]** or **[edit interfaces *interface-name* ether-options]** hierarchy level.

SEE ALSO

Ethernet Interfaces User Guide for Routing Devices

For information about configuring aggregated devices, see the *Junos OS Administration Library*.

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 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/1 {
  gigether-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 Aggregate Ethernet interface bundles](#)

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](#)

[Periodic Packet Management](#)

Understanding Local Link Bias


Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF using a different member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is enabled is illustrated in [Figure 7 on page 274](#).

Figure 7: Egress Traffic Flow with Local Link Bias



When local link bias is disabled, egress traffic exiting a Virtual Chassis or VCF on a LAG bundle can be forwarded out of any member link in the LAG bundle. Traffic forwarding decisions are made by an internal algorithm that attempts to load-balance traffic between the member links in the bundle. VCP bandwidth is frequently consumed by egress traffic when local link bias is disabled because the egress traffic traverses the VCPs to reach the destination egress member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is disabled is illustrated in [Figure 8 on page 274](#).

Figure 8: Egress Traffic Flow without Local Link Bias



Starting in Junos OS Release 14.1X53-D25, local link bias can be enabled globally for all LAG bundles in a Virtual Chassis or VCF, or individually per LAG bundle in a Virtual Chassis. In prior Junos OS releases, local link bias could be enabled individually per LAG bundle only.

A Virtual Chassis or VCF that has multiple LAG bundles can contain bundles that have and have not enabled local link bias. Local link bias only impacts the forwarding of unicast traffic exiting a Virtual Chassis or VCF; ingress traffic handling is not impacted by the local link bias setting. Egress multicast, unknown unicast, and broadcast traffic exiting a Virtual Chassis or VCF over a LAG bundle is not impacted by the local link bias setting and is always load-balanced among the member links. Local link bias is disabled, by default.

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

Configuring Local Link Bias

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

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

Local link bias can be enabled or disabled globally or per LAG bundle on a Virtual Chassis or VCF. In cases where local link bias is enabled at both the global and per LAG bundle levels, the per LAG bundle configuration takes precedence. For instance, if local link bias is enabled globally but disabled on a LAG bundle named **ae1**, local link bias is disabled on the LAG bundle named **ae1**.

To enable local link bias on a LAG bundle:

```
[edit]
user@switch# set interface aex aggregated-ether-options local-bias
```

where **aex** is the name of the aggregated Ethernet link bundle.

For instance, to enable local link bias on aggregated Ethernet interface **ae0**:

```
[edit]
user@switch# set interface ae0 aggregated-ether-options local-bias
```

Understanding Local Minimum Links

IN THIS SECTION

- [Configuring Local Minimum Links | 278](#)
- [Local Minimum Links Effect on LAG Minimum Links | 279](#)
- [Local Minimum Links and Local Link Bias | 279](#)

NOTE: When describing the local minimum links feature, *member links* are links that are part of an aggregated Ethernet bundle (LAG), *member switches* are chassis that are members in a Virtual Chassis or Virtual Chassis Fabric (VCF), and *local member links* (or simply *local links*) are member links of the same LAG that are local to a particular Virtual Chassis or VCF member switch.

A link aggregation group (LAG) can include member links on different chassis, and multiple local member links on member switches in a Virtual Chassis or VCF. If member links in the LAG fail, the LAG continues to carry traffic over the remaining member links that are still active. When multiple member links are local to one chassis and one or more of those links fail, LAG traffic coming into that chassis will be redistributed over the remaining local links. However, the remaining active local links can suffer traffic loss if the failed links result in sufficiently reduced total bandwidth through the chassis.

Introduced in Junos OS Release 14.1X53-D40, the local minimum links feature helps avoid traffic loss due to asymmetric bandwidth on LAG forwarding paths through a Virtual Chassis or VCF member switch when one or more local member links have failed.

NOTE: The local minimum links feature is supported on Virtual Chassis or VCFs with QFX5100 member switches only.

Based on a user-configured threshold value, when one or more member links fail, this feature marks any remaining active local links as “down,” forcing LAG traffic to be redistributed only through member links on *other* chassis. To enable this feature on a particular aggregated Ethernet interface (aex), you set the **local-minimum-links-threshold** configuration statement with a threshold value that represents the percentage of local member links that must be up on a chassis for *any* local member links on that chassis to continue to be active in the aggregated Ethernet bundle.

The configured threshold value:

- Applies to a specified aggregated Ethernet interface.
- Applies to any chassis that has links in the specified aggregated Ethernet bundle.
- Represents a percentage of active local member links out of the total number of local member links for the chassis.

When the local minimum links feature is enabled for a LAG, if one or more member links on a chassis fail, the feature compares the percentage of local member links that are still up to the threshold. If the percentage of “up” links is less than the threshold, the feature forces down the remaining active local links, and no traffic for the aggregated Ethernet interface will be forwarded through the member links on that chassis. If the percentage of links that are “up” is greater than or equal to the threshold, the status of the active links remains unchanged, and LAG traffic will continue to be distributed over available member links on that chassis.

For example, consider a member switch in a Virtual Chassis Fabric that has four links that are active member links of a LAG, and the local minimum links feature is enabled with the threshold set to 60:

- If one member link goes down, 75 percent (three out of four) of the links are still up, which is greater than the threshold (60 percent), so the remaining links stay up.
- If two member links go down, only 50 percent (two out of four) of the links are “up”, so the local minimum links feature forces the remaining two active links “down.” The same is true if three member links fail, the remaining link is forced down as well.

The local minimum links feature tracks whether links are down because the link failed or the link was forced down, as well as when active, failed, or forced-down member links are added or removed. As a result, the feature can respond dynamically when:

- Failed local member links come back up.
- You change the configured threshold value, or you disable the local minimum links feature.
- Adding or removing local member links changes the total number of local member links, or changes the ratio of “up” links to total local member links as compared to the threshold.

For example, if a failed member link causes all local member links to be forced down, then that link comes back up and brings the percentage of “up” links above the current threshold, the system adjusts the status of the forced-down links to mark them up again as well.

You should enable this feature only if your system closely manages ingress and egress traffic forwarding paths on LAGs for individual chassis in a Virtual Chassis and VCFs, especially where local link bias is also enabled.

Configuring Local Minimum Links

The local minimum links feature is disabled by default. To enable this feature for a LAG bundle (which then applies to any chassis that has local member links in the LAG), simply configure a threshold value for the LAG interface, as follows:

```
[edit interfaces]
user@switch# set aggregated-ether-options aex local-minimum-links-threshold threshold-value
```

To update the threshold value, use the same command with the new threshold value.

To disable the local minimum links feature, delete the **local-minimum-links-threshold** statement from the configuration. Any links that were forced down by this feature are automatically brought up again within a few seconds.

Local Minimum Links Effect on LAG Minimum Links

The per-chassis local minimum links threshold is similar to the [minimum-links](#) setting for a LAG bundle, which configures the minimum number of member links in the bundle that should be up for the aggregated Ethernet interface as a whole to be considered “up.” (See [“Configuring Link Aggregation” on page 281](#) for details.) Local member links that fail or are forced down by the local minimum links feature contribute to the count of “up” links for the LAG as a whole. As a result, this feature can cause the entire LAG to be brought down if enough local links are forced down. Enabling and configuring the local minimum links feature is independent of LAG minimum links configuration, but you should carefully consider the combined potential effect on the LAG as a whole when configuring both features.

Local Minimum Links and Local Link Bias

The local minimum links and local link bias features operate independently, but can influence each other’s traffic forwarding results. For example, when local link bias is enabled and would otherwise favor forwarding traffic out of local links in the aggregated Ethernet bundle, but those links are down because the local minimum links threshold is not currently met, outgoing traffic will be redirected through the VCPs to other Virtual Chassis or VCF member switches for forwarding. In that case, unanticipated increased VCP traffic can impact Virtual Chassis or VCF performance.

See [“Understanding Local Link Bias” on page 274](#) for details on the local link bias feature.

SEE ALSO

| [local-minimum-links-threshold](#) | [617](#)

Troubleshooting an Aggregated Ethernet Interface

IN THIS SECTION

- [Show Interfaces Command Shows the LAG is Down](#) | [280](#)
- [Logical Interface Statistics Do Not Reflect All Traffic](#) | [280](#)
- [IPv6 Interface Traffic Statistics Are Not Supported](#) | [280](#)
- [SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0](#) | [280](#)

Troubleshooting issues for aggregated Ethernet interfaces:

Show Interfaces Command Shows the LAG is Down

Problem

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

Solution

Check the following:

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

Logical Interface Statistics Do Not Reflect All Traffic

Problem

Description: The traffic statistics for a logical interface do not include all of the traffic.

Solution

Traffic statistics fields for logical interfaces in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic. You can view the statistics for all traffic only per physical interface.

IPv6 Interface Traffic Statistics Are Not Supported

Problem

Description: The IPv6 transit statistics in the **show interfaces** command display all 0 values.

Solution

EX Series switches do not support the collection and reporting of IPv6 transit statistics.

SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0

Problem

Description: The values for the SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are always 0.

Solution

The SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are not supported for aggregated Ethernet interfaces on EX Series switches.

SEE ALSO

| *Verifying the Status of a LAG Interface*

Configuring Link Aggregation

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

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

NOTE: On QFX5100, QFX5200, EX4600, QFX10002, and QFX10008 standalone switches and on QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Load balancing will not work if you configure link speeds that are not supported. (Platform support depends on the Junos OS release in your installation.)

1. [Creating an Aggregated Ethernet Interface | 282](#)
2. [Configuring the VLAN Name and VLAN ID Number | 283](#)
3. [Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

Creating an Aggregated Ethernet Interface

To create an aggregated Ethernet interface:

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

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

For example, to specify 5:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count 5
```

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

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

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

For example, to specify 5:

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

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

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

For example, to specify 10g:

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

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

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

Configuring the VLAN Name and VLAN ID Number

NOTE: VLANs are not supported on OCX Series switches.

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

For example, 100.

NOTE: When you add or remove a vlan from a LAG interface, the interface goes down and comes back (flaps). The flapping happens when a low speed SFP is plugged into a relatively high speed port. To avoid flapping, configure the port speed to match the speed of the SFP.

Configuring Aggregated Ethernet LACP (CLI Procedure)

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

LACP was designed to achieve the following:

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

NOTE: You can also configure LACP link protection on aggregated Ethernet interfaces. For information, see [“Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches” on page 313](#).

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

Before you configure LACP for EX Series, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See *Configuring Aggregated Ethernet Links (CLI Procedure)*

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

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

To configure LACP:

1. Enable the LACP mode:

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

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

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

NOTE: LACP decides active and back up state of links. When configuring LACP, state of the backup link should not be configured manually as down. The following command is not supported if LACP is configured: `set interfaces ae0 aggregated-ether-options link-protection backup-state down`

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

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

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

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

3. (Optional) A link without Link Access Control Protocol (LACP) configuration remains down and cannot be accessed by the provider edge (PE) devices in the topology. Configure the **force-up** feature in LACP on a PE device for which you need connectivity.

```
[edit interfaces]
```

```
user@switch# set interfaces aex aggregated-ether-options lacp force-up
```

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

SEE ALSO

Configuring Aggregated Ethernet Links (CLI Procedure)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

Configuring Aggregated Ethernet Interfaces (J-Web Procedure)

[Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 321](#)

[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)

Verifying the Status of a LAG Interface

SEE ALSO

[Understanding Interface Naming Conventions | 38](#)

Configuring an FCoE LAG

[Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch | 299](#)

Verifying the Status of a LAG Interface

[Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets | 320](#)

[show lacp statistics interfaces \(View\) | 1308](#)

Configuring Aggregated Ethernet Link Protection

IN THIS SECTION

- [Configuring Link Protection for Aggregated Ethernet Interfaces | 286](#)
- [Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces | 287](#)
- [Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link | 287](#)
- [Disabling Link Protection for Aggregated Ethernet Interfaces | 288](#)

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](#) | [606](#)

aggregated-ether-options

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

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

To configure a primary link and a backup link:

1. Configure the primary logical interface.

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

2. Configure the backup logical interface.

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

SEE ALSO

802.3ad

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

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

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

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

SEE ALSO

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

Disabling Link Protection for Aggregated Ethernet Interfaces

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

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

SEE ALSO

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

Configuring Aggregated Ethernet 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*

Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface

If subscribers are frequently logging in and logging out of your network, you can configure the system to periodically rebalance the links based on a specific time and interval.

To configure periodic rebalancing:

1. Access the aggregated Ethernet interface for which you want to configure periodic rebalancing.

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

2. Configure the rebalancing parameters for the interface, including the time and the interval between rebalancing actions.

```
[edit interfaces aenumber aggregated-ether-options]
user@host# rebalance-periodic time hour:minute <interval hours>
```

SEE ALSO

Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface

Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces

Distribution of Demux Subscribers in an Aggregated Ethernet Interface

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

IN THIS SECTION

- [Requirements | 292](#)
- [Overview and Topology | 292](#)
- [Configuration | 294](#)
- [Verification | 297](#)
- [Troubleshooting | 298](#)

EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your EX Series switch model.

This example describes how to configure uplink LAGs to connect a Virtual Chassis access switch to a Virtual Chassis distribution switch:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches
- Two EX4200-48P switches
- Two EX4200-24F switches
- Four XFP uplink modules

Before you configure the LAGs, be sure you have:

- Configured the Virtual Chassis switches. See *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*.
- Configured the uplink ports on the switches as trunk ports. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*.

Overview and Topology

For maximum speed and resiliency, you can combine uplinks between an access switch and a distribution switch into LAGs. Using LAGs can be particularly effective when connecting a multimember Virtual Chassis access switch to a multimember Virtual Chassis distribution switch.

The Virtual Chassis access switch in this example is composed of two member switches. Each member switch has an uplink module with two 10-Gigabit Ethernet ports. These ports are configured as trunk ports, connecting the access switch with the distribution switch.

Configuring the uplinks as LAGs has the following advantages:

- Link Aggregation Control Protocol (LACP) can optionally be configured for link negotiation.
- It doubles the speed of each uplink from 10 Gbps to 20 Gbps.
- If one physical port is lost for any reason (a cable is unplugged or a switch port fails, or one member switch is unavailable), the logical port transparently continues to function over the remaining physical port.

The topology used in this example consists of one Virtual Chassis access switch and one Virtual Chassis distribution switch. The access switch is composed of two EX4200-48P switches (SWA-0 and SWA-1), interconnected to each other with their Virtual Chassis ports (VCPs) as member switches of Host-A. The distribution switch is composed of two EX4200-24F switches (SWD-0 and SWD-1), interconnected with their VCPs as member switches of Host-D.

Each member of the access switch has an uplink module installed. Each uplink module has two ports. The uplinks are configured to act as trunk ports, connecting the access switch with the distribution switch. One uplink port from SWA-0 and one uplink port from SWA-1 are combined as LAG ae0 to SWD-0. This

link is used for one VLAN. The remaining uplink ports from SWA-0 and from SWA-1 are combined as a second LAG connection (**ae1**) to SWD-1. LAG **ae1** is used for another VLAN.

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

Figure 9: Topology for LAGs Connecting an EX4200 Virtual Chassis Access Switch to an EX4200 Virtual Chassis Distribution Switch

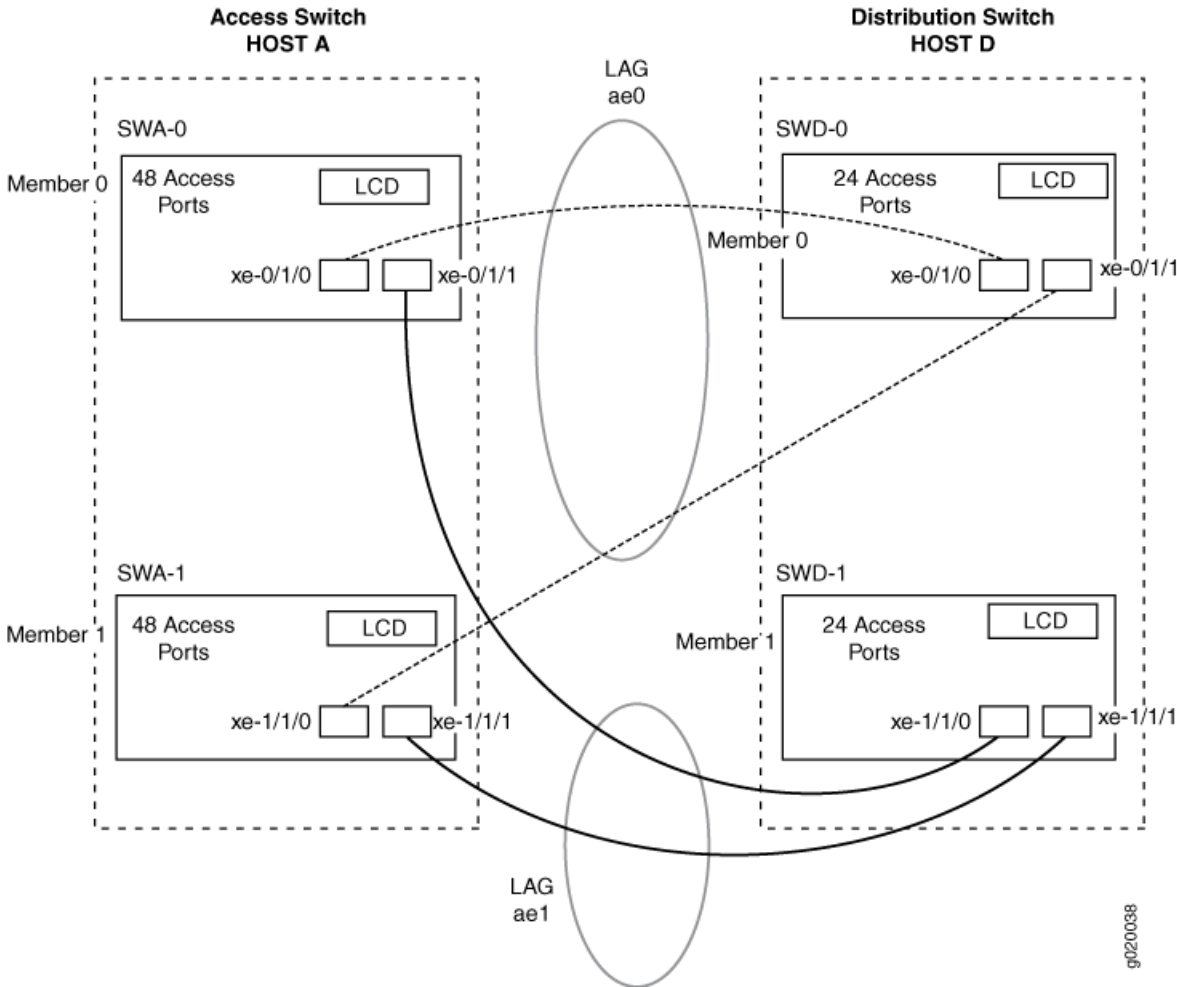


Table 53 on page 294 details the topology used in this configuration example.

Table 53: Components of the Topology for Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWA-0	Host-A Access switch VCID 1	EX4200-48P switch	One XFP uplink module	0	xe-0/1/0 to SWD-0 xe-0/1/1 to SWD-1
SWA-1	Host-A Access switch VCID 1	EX4200-48P switch	One XFP uplink module	1	xe-1/1/0 to SWD-0 xe-1/1/1 to SWD-1
SWD-0	Host-D Distribution switch VCID 4	EX4200 L-24F switch	One XFP uplink module	0	xe-0/1/0 to SWA-0 xe-0/1/1 to SWA-1
SWD-1	Host-D Distribution switch VCID 4	EX4200 L-24F switch	One XFP uplink module	1	xe-1/1/0 to SWA-0 xe-1/1/1 to SWA-1

Configuration

To configure two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch:

CLI Quick Configuration

To quickly configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch, copy the following commands and paste them into the switch terminal window:

[edit]

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options minimum-links 1
```

```

set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family inet address 192.0.2.0/25
set interfaces ae1 unit 0 family inet address 192.0.2.128/25
set interfaces xe-0/1/0 ether-options 802.3ad ae0
set interfaces xe-1/1/0 ether-options 802.3ad ae0
set interfaces xe-0/1/1 ether-options 802.3ad ae1
set interfaces xe-1/1/1 ether-options 802.3ad ae1

```

Step-by-Step Procedure

To configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch:

1. Specify the number of LAGs to be created on the chassis:

```

[edit chassis]
user@Host-A# set aggregated-devices ethernet device-count 2

```

2. Specify the number of links that need to be present for the **ae0** LAG interface to be **up**:

```

[edit interfaces]
user@Host-A# set ae0 aggregated-ether-options minimum-links 1

```

3. Specify the number of links that need to be present for the **ae1** LAG interface to be **up**:

```

[edit interfaces]
user@Host-A# set ae1 aggregated-ether-options minimum-links 1

```

4. Specify the media speed of the **ae0** link:

```

[edit interfaces]
user@Host-A# set ae0 aggregated-ether-options link-speed 10g

```

5. Specify the media speed of the **ae1** link:

```

[edit interfaces]
user@Host-A# set ae1 aggregated-ether-options link-speed 10g

```

6. Specify the interface ID of the uplinks to be included in LAG **ae0**:

```

[edit interfaces]
user@Host-A# set xe-0/1/0 ether-options 802.3ad ae0
user@Host-A# set xe-1/1/0 ether-options 802.3ad ae0

```

7. Specify the interface ID of the uplinks to be included in LAG **ae1**:

```
[edit interfaces]
user@Host-A# set xe-0/1/1 ether-options 802.3ad ae1
user@Host-A# set xe-1/1/1 ether-options 802.3ad ae1
```

8. Specify that LAG **ae0** belongs to the subnet for the employee broadcast domain:

```
[edit interfaces]
user@Host-A# set ae0 unit 0 family inet address 192.0.2.0/25
```

9. Specify that LAG **ae1** belongs to the subnet for the guest broadcast domain:

```
[edit interfaces]
user@Host-A# set ae1 unit 0 family inet address 192.0.2.128/25
```

Results

Display the results of the configuration:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  ae0 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
      family inet {
        address 192.0.2.0/25;
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
```



```

        family inet {
            address 192.0.2.128/25;
        }
    }
    xe-0/1/0 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-1/1/0 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/1/1 {
        ether-options {
            802.3ad ae1;
        }
    }
    xe-1/1/1 {
        ether-options {
            802.3ad ae1;
        }
    }
}

```

Verification

IN THIS SECTION

- [Verifying That LAG ae0 Has Been Created | 297](#)
- [Verifying That LAG ae1 Has Been Created | 298](#)

To verify that switching is operational and two LAGs have been created, perform these tasks:

Verifying That LAG ae0 Has Been Created

Purpose

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

Action

show interfaces ae0 terse

Interface	Admin	Link	Proto	Local	Remote
ae0	up	up			
ae0.0	up	up	inet	192.0.2.0/25	

Meaning

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

Verifying That LAG ae1 Has Been Created**Purpose**

Verify that LAG **ae1** has been created on the switch

Action

show interfaces ae1 terse

Interface	Admin	Link	Proto	Local	Remote
ae1	up	down			
ae1.0	up	down	inet	192.0.2.128/25	

Meaning

The output shows that the **ae1** link is down.

Troubleshooting**Troubleshooting a LAG That Is Down****Problem**

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

Solution

Check the following:

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

SEE ALSO

Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet

Example: Connecting an EX Series Access Switch to a Distribution Switch.

Virtual Chassis Cabling Configuration Examples for EX4200 Switches

Installing an Uplink Module in an EX4200 Switch

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

IN THIS SECTION

- [Requirements | 299](#)
- [Overview and Topology | 300](#)
- [Configuration | 300](#)
- [Verification | 303](#)
- [Troubleshooting | 304](#)

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

Requirements

This example uses the following software and hardware components:

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

Overview and Topology

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

Configuring the Ethernet interfaces as LAGs has the following advantages:

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

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

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




Table 54 on page 300 details the topology used in this configuration example.

Table 54: Components of the Topology for Configuring a LAG Between a Switch and an Aggregation Switch

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

Configuration

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

CLI Quick Configuration

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

NOTE: If you are configuring a LAG using Enhanced Layer 2 Software—for example, on the EX4600, QFX5100, or QFX10002 switch—use the **interface-mode** statement instead of the **port-mode** statement. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

```
[edit]
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family ethernet-switching vlan members green
set interfaces xe-0/0/2 ether-options 802.3ad ae0
set interfaces xe-0/0/3 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
```

Step-by-Step Procedure

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

1. Specify the number of LAGs to be created on the switch:

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

2. Specify the number of links that need to be present for the **ae0** LAG interface to be **up**:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links 1
```

3. Specify the media speed of the **ae0** link:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed 10g
```

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

```
[edit interfaces]
```

```
user@switch# set interfaces xe-0/0/2 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set interfaces xe-0/0/3 ether-options 802.3ad ae0
```

5. Assign a port mode of trunk to the **ae0** link:

NOTE: If you are configuring a LAG using Enhanced Layer 2 Software—for example, on the EX4600, QFX5100, or QFX10002 switch—use the **interface-mode** statement instead of the **port-mode** statement. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

or

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

6. Assign the LAG to a VLAN:

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members green vlan-id 200
```

7. (Optional): Designate one side of the LAG as active for LACP:

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

8. (Optional): Designate the interval and speed at which the interfaces send LACP packets:

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

Results

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

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
```

```

    }
  }
}
green {
  vlan-id 200;
}
}
interfaces {
  ae0 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members green;
        }
      }
    }
  }
  xe-0/0/2 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/3 {
    ether-options {
      802.3ad ae0;
    }
  }
}

```

Verification

IN THIS SECTION

- [Verifying That LAG ae0.0 Has Been Created | 304](#)
- [Verifying That LAG ae0 Has Been Created | 304](#)

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

Verifying That LAG ae0.0 Has Been Created

Purpose

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

Action

show interfaces ae0 terse

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

Meaning

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

Verifying That LAG ae0 Has Been Created

Purpose

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

Action

show interfaces ae0 terse

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

Meaning

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

Troubleshooting

Troubleshooting a LAG That Is Down

Problem

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

Solution

Check the following:

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

SEE ALSO

Verifying the Status of a LAG Interface

[show lacp statistics interfaces \(View\) | 1308](#)

Configuring Aggregated Ethernet LACP

IN THIS SECTION

- [Configuring the LACP Interval | 307](#)
- [Configuring LACP Link Protection | 307](#)
- [Configuring LACP System Priority | 309](#)
- [Configuring LACP System Identifier | 309](#)
- [Configuring LACP administrative Key | 310](#)
- [Configuring LACP Port Priority | 310](#)
- [Tracing LACP Operations | 311](#)
- [LACP Limitations | 311](#)
- [Example: Configuring Aggregated Ethernet LACP | 311](#)

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.

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.

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;
}
```

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.

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: Source address filtering does not work when LACP is enabled.

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*.

Configuring LACP Link Protection

NOTE: When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.

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 **[ether-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: LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.

To enable LACP link protection for an aggregated Ethernet interfaces, 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 you to use LACP on both ends of the aggregator, when you connect an aggregated Ethernet interface with two member interfaces 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-SFPP 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;
```

The user-defined system identifier in LACP enables two ports from two separate devices to act as though they were part of the same aggregate group.

The system identifier is a 48-bit (6-byte) globally unique field. It is used in combination with a 16-bit system-priority value, which results in a unique LACP system identifier.

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* ether-options 802.3ad aeX lacp]** or **[edit interfaces *interface-name* ether-options 802.3ad aeX lacp]** hierarchy levels:

```
[edit interfaces interface-name ether-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 are 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: Port aggregation selection (discussed above) 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.

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;
}
```

You can specify the following flags in the **protocols lacp traceoptions** statement:

- **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

For general information about tracing, see the tracing and logging information in the *Junos OS Administration Library*.

LACP Limitations

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.

Example: Configuring Aggregated Ethernet LACP

Configure aggregated Ethernet LACP over a VLAN-tagged interface:

LACP with VLAN-Tagged Aggregated Ethernet

```

[edit interfaces]
ge--1/1/1 {
  ether-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;
        }
      }
    }
  }
}

```

Configure aggregated Ethernet LACP over an untagged interface:

LACP with Untagged Aggregated Ethernet

```

[edit interfaces]
ge-1/1/1 {
  ether-options-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

[lACP](#) | [600](#)

[link-protection](#) | [606](#)

[traceoptions](#)

Ethernet Interfaces User Guide for Routing Devices

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

IN THIS SECTION

- [Configuring LACP Link Protection for a Single Link at the Global Level](#) | [315](#)
- [Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level](#) | [315](#)
- [Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface](#) | [316](#)

You can configure LACP link protection and system priority at the global level on the switch or for a specific aggregated Ethernet interface. When using LACP link protection to protect a single link in the aggregated ethernet bundle, you configure only two member links for an aggregated Ethernet interface: one active and one standby. LACP link protection ensures that only one link—the link with the higher priority—is used for traffic. The other link is forced to stay in a *waiting* state.

When using LACP link protection to protect multiple links in an aggregated ethernet bundle, you configure links into primary and backup subgroups. A link protection subgroup is a collection of ethernet links within the aggregated ethernet bundle. When you use link protection subgroups, you configure a primary subgroup and a backup subgroup. The configuration process includes assigning member links to each subgroup. When the configuration process is complete, the primary subgroup is used to forward traffic until a switchover event, such as a link failure, occurs and causes the backup subgroup to assume control of traffic that was travelling on the links in the primary subgroup within the bundle.

By default LACP link protection reverts to a higher-priority (lower-numbered) link when the higher-priority link becomes operational or when a higher-priority link is added to the aggregated Ethernet bundle. For priority purposes, LACP link protection treats subgroups like links. You can suppress link calculation by adding the **non-revertive** statement to the link protection configuration. In nonrevertive mode, when a link is active in sending and receiving LACP packets, adding a higher-priority link to the bundle does not change the status of the currently active link. It remains active.

If LACP link configuration is specified to be nonrevertive at the global **[edit chassis]** hierarchy level, you can specify the **revertive** statement in the LACP link protection configuration at the aggregated Ethernet interface level to override the nonrevertive setting for the interface. In revertive mode, adding a higher-priority link to the aggregated Ethernet bundle results in LACP recalculating the priority and switching the status from the currently active link to the newly added, higher-priority link.

NOTE: When LACP link protection is enabled on both local and remote sides of the link, both sides must use the same mode (either revertive or nonrevertive).

Configuring LACP link configuration at the aggregated Ethernet level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

Before you configure LACP link protection, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). For EX Series, see *Configuring Aggregated Ethernet Links (CLI Procedure)*.

- Configured LACP for the interface. For Ex Series, see [“Configuring Aggregated Ethernet LACP \(CLI Procedure\)” on page 283](#).

You can configure LACP link protection for all aggregated Ethernet interfaces on the switch by enabling it at the global level on the switch or configure it for a specific aggregated Ethernet interface by enabling it on that interface.

Configuring LACP Link Protection for a Single Link at the Global Level

To configure LACP link protection for aggregated Ethernet interfaces at the global level:

1. Enable LACP link protection on the switch:

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set link-protection
```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interfaces to be in nonrevertive mode:

NOTE: LACP link protection is in revertive mode by default.

```
[edit chassis aggregated-devices ethernet lacp link-protection]
user@switch# set non-revertive
```

3. (Optional) To configure LACP system priority for the aggregated Ethernet interfaces:

```
[edit chassis aggregated-devices ethernet lacp]
user@switch# set system-priority
```

Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level

To enable LACP link protection for a specific aggregated Ethernet interface:

1. Enable LACP link protection for the interface:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interface to be in revertive or nonrevertive mode:

- To specify revertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set revertive
```

- To specify nonrevertive mode:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set non-revertive
```

3. (Optional) To configure LACP system priority for an aggregated Ethernet interface:

```
[edit interfaces aeX aggregated-ether-options lacp link-protection]
user@switch# set system-priority
```

4. (Optional) To configure LACP port priority for an aggregated Ethernet interface:

```
[edit interfaces ge-fpc/pic/port ether-options 802.3ad lacp]
user@switch# set port-priority
```

Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface

You can configure link protection subgroup bundles to provide link protection for multiple links in an aggregated ethernet bundle.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a LAG bundle, instead of providing protection to a single link in the aggregated ethernet bundle only. You can, for instance, configure a primary subgroup with three member links and a backup subgroup with three different member links and use the backup subgroup to provide link protection for the primary subgroup.

To configure link protection using subgroups:

1. Configure the primary link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection-sub-group group-name primary
```

For instance, to create a primary link protection subgroup named **subgroup-primary** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection-sub-group subgroup-primary primary
```

2. Configure the backup link protection subgroup in the aggregated ethernet interface:

```
[edit interfaces aeX aggregated-ether-options]
```

```
user@switch# set link-protection-sub-group group-name backup
```

For instance, to create a backup link protection subgroup named **subgroup-backup** for interface **ae0**:

```
[edit interfaces ae0 aggregated-ether-options]
```

```
user@switch# set link-protection-sub-group subgroup-backup backup
```

NOTE: You can create one primary and one backup link protection subgroup per aggregated ethernet interface.

3. Attach interfaces to the link protection subgroups:

```
[edit interfaces interface-name ether-options 802.3ad]
```

```
user@switch# set link-protection-sub-group group-name
```

NOTE: The primary and backup link protection subgroups must contain the same number of interfaces. For instance, if the primary link protection subgroup contains three interfaces, the backup link protection subgroup must also contain three interfaces.

For instance, to configure interfaces **ge-0/0/0** and **ge-0/0/1** into link protection subgroup **subgroup-primary** and interfaces **ge-0/0/2** and **ge-0/0/3** into link protection subgroup **subgroup-backup**:

```
[edit interfaces ge-0/0/0 ether-options 802.3ad]
```

```
user@switch# set link-protection-sub-group subgroup-primary
```

```
[edit interfaces ge-0/0/1 ether-options 802.3ad]
```

```
user@switch# set link-protection-sub-group subgroup-primary
```

```
[edit interfaces ge-0/0/2 ether-options 802.3ad]
```

```
user@switch# set link-protection-sub-group subgroup-backup
```

```
[edit interfaces ge-0/0/3 ether-options 802.3ad]
```

```
user@switch# set link-protection-sub-group subgroup-backup
```

4. (Optional) Configure the port priority for link protection:

```
[edit interfaces interface-name ether-options 802.3ad]
```

```
user@switch# set port-priority priority
```

The port priority is used to select the active link.

5. Enable link protection

To enable link protection at the LAG level:

```
[edit interfaces aeX aggregated-ether-options]
user@switch# set link-protection
```

NOTE: ACX Series routers do not support static link protection.

To enable link protection at the LACP level:

```
[edit interfaces aeX aggregated-ether-options lacp]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LAG level:

```
[edit interfaces ae0 aggregated-ether-options]
user@switch# set link-protection
```

For instance, to enable link protection on **ae0** at the LACP level:

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

NOTE: The LACP decides active and back up state of links. When configuring LACP, the state of the backup link should not be configured manually as down. The following command is not supported if LACP is configured:
set interfaces ae0 aggregated-ether-options link-protection backup-state down

SEE ALSO

| [lacp \(Aggregated Ethernet\)](#) | 600

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

Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

Verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

- 1. [Verifying the LACP Setup | 320](#)
- 2. [Verifying That LACP Packets Are Being Exchanged | 320](#)

Verifying the LACP Setup

Purpose

Verify that the LACP has been set up correctly.

Action

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

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

Aggregated interface: ae0									
LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/1/0	Actor	No	Yes	No	No	No	Yes	Fast	Active
xe-0/1/0	Partner	No	Yes	No	No	No	Yes	Fast	Passive
LACP protocol:	Receive State		Transmit State					Mux	State
xe-0/1/0	Defaulted		Fast periodic					Detached	

Meaning

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

Verifying That LACP Packets Are Being Exchanged

Purpose

Verify that LACP packets are being exchanged between interfaces.

Action

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

show lacp statistics interfaces ae0


```

Aggregated interface: ae0
  LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
    xe-0/0/2            1352        2035           0                0
    xe-0/0/3            1352        2056           0                0
  
```

Meaning

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

SEE ALSO

Verifying the Status of a LAG Interface

[show lacp statistics interfaces \(View\) | 1308](#)

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

IN THIS SECTION

- [Requirements | 322](#)
- [Overview and Topology | 322](#)
- [Configuring LACP for the LAGs on the Virtual Chassis Access Switch | 323](#)
- [Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch | 324](#)
- [Verification | 325](#)
- [Troubleshooting | 327](#)

EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. EX Series switches allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in [“Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch”](#) on page 291:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches
- Two EX4200-48P switches
- Two EX4200-24F switches
- Four EX Series XFP uplink modules

Before you configure LACP, be sure you have:

- Set up the Virtual Chassis switches. See *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*.
- Configured the uplink ports on the switches as trunk ports. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*.
- Configured the LAGs. See [“Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch”](#) on page 291.

Overview and Topology

This example assumes that you are familiar with [“Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch”](#) on page 291. The topology in this example is exactly the same as the topology in that other example. This example shows how to use LACP to enhance the LAG functionality.

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

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

By default, the actor and partner send LACP packets every second.

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

Configuring LACP for the LAGs on the Virtual Chassis Access Switch

To configure LACP for the access switch LAGs, perform these tasks:

CLI Quick Configuration

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

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
set interfaces ae1 aggregated-ether-options lacp active periodic fast
```

Step-by-Step Procedure

To configure LACP for Host-A LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-A#set ae0 aggregated-ether-options lacp active periodic fast
user@Host-A#set ae1 aggregated-ether-options lacp active periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
user@Host-A# show
ae0 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
```

Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch

To configure LACP for the two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch, perform these tasks:

CLI Quick Configuration

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

```
[edit interfaces]
set ae0 aggregated-ether-options lacp passive periodic fast
set ae1 aggregated-ether-options lacp passive periodic fast
```

Step-by-Step Procedure

To configure LACP for Host D LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-D#set ae0 aggregated-ether-options lacp passive periodic fast
user@Host-D#set ae1 aggregated-ether-options lacp passive periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
user@Host-D# show
ae0 {
  aggregated-ether-options {
    lacp {
      passive;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      passive
      periodic fast;
    }
  }
}
```

```
}
```

Verification

IN THIS SECTION

- [Verifying the LACP Settings | 325](#)
- [Verifying That the LACP Packets Are Being Exchanged | 326](#)

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

Verifying the LACP Settings

Purpose

Verify that LACP has been set up correctly.

Action

Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@Host-A> show lacp interfaces xe-0/1/0
```

Aggregated interface: ae0									
LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/1/0	Actor	No	Yes	No	No	No	Yes	Fast	Active
xe-0/1/0	Partner	No	Yes	No	No	No	Yes	Fast	Passive
LACP protocol:	Receive State			Transmit State			Mux State		
xe-0/1/0	Defaulted			Fast periodic			Detached		

Meaning

The output indicates that LACP has been set up correctly and is active at one end.

Verifying That the LACP Packets Are Being Exchanged

Purpose

Verify that LACP packets are being exchanged.

Action

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

user@Host-A> **show interfaces ae0 statistics**

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

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

Meaning

The output here shows that the link is down and that no protocol data units (PDUs) are being exchanged.

Troubleshooting

IN THIS SECTION

- [Troubleshooting a Nonworking LACP Link | 327](#)

To troubleshoot a nonworking LACP link, perform these tasks:

Troubleshooting a Nonworking LACP Link

Problem

The LACP link is not working.

Solution

Check the following:

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

SEE ALSO

Example: Connecting an EX Series Access Switch to a Distribution Switch

Virtual Chassis Cabling Configuration Examples for EX4200 Switches

Installing an Uplink Module in an EX4200 Switch

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

IN THIS SECTION

- Requirements | 328
- Overview and Topology | 329
- Configuring LACP for the LAG on the QFX Series | 329
- Verification | 330
- Troubleshooting | 332

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

This example describes how to overlay LACP on the LAG configurations that were created in [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch” on page 299](#):

Requirements

This example uses the following software and hardware components:

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

Before you configure LACP, be sure you have:

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

Overview and Topology

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

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

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

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

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

Configuring LACP for the LAG on the QFX Series

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

CLI Quick Configuration

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

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

Step-by-Step Procedure

To configure LACP for LAG **ae0** :

1. Specify the aggregated Ethernet options for the LAG:

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

Results

Display the results of the configuration:

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

Verification

IN THIS SECTION

- [Verifying the LACP Settings | 330](#)
- [Verifying That the LACP Packets Are Being Exchanged | 331](#)

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

Verifying the LACP Settings

Purpose

Verify that LACP has been set up correctly.

Action

Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

user@switch> show lacp interfaces xe-0/0/2

Aggregated interface: ae0									
LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/2	Actor	No	Yes	No	No	No	Yes	Fast	Active
xe-0/0/2	Partner	No	Yes	No	No	No	Yes	Fast	Passive

LACP protocol:	Receive State	Transmit State	Mux State
xe-0/0/2	Defaulted	Fast periodic	Detached

Meaning

The output indicates that LACP has been set up correctly and is active at one end.

Verifying That the LACP Packets Are Being Exchanged

Purpose

Verify that LACP packets are being exchanged.

Action

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

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

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

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
  Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
  Statistics          Packets          pps          Bytes          bps
  Bundle:
    Input :              0              0              0              0
    Output:              0              0              0              0
  Protocol inet
    Flags: None
```

```
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.10.10/8, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning

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

Troubleshooting

IN THIS SECTION

- [Troubleshooting a Nonworking LACP Link | 332](#)

To troubleshoot a nonworking LACP link, perform these tasks:

Troubleshooting a Nonworking LACP Link

Problem

The LACP link is not working.

Solution

Check the following:

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

SEE ALSO

Verifying the Status of a LAG Interface

[Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch | 299](#)

Example: Configuring an FCoE LAG on a Redundant Server Node Group

[show lacp statistics interfaces \(View\) | 1308](#)

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](#) | 561

detection-time

transmit-interval

Configuring Micro BFD Sessions for LAG

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.

NOTE: Starting in Junos OS Evolved Release 20.1R1, independent micro Bidirectional Forwarding Detection (BFD) sessions are enabled on a per member link basis of a Link Aggregation Group (LAG) bundle.

To enable failure detection for aggregated Ethernet interfaces:

1. Include the following statement in the configuration at the **[edit interfaces *aex* aggregated-ether-options]** hierarchy level:

```
bfd-liveness-detection
```

2. Configure the authentication criteria of the BFD session for LAG.

To specify the authentication criteria, include the **authentication** statement:

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
}
```

- Specify the algorithm to be used to authenticate the BFD session. You can use one of the following algorithms for authentication:
 - keyed-md5
 - keyed-sha-1
 - meticulous-keyed-md5
 - meticulous-keyed-sha-1
 - simple-password

- To configure the key chain, specify the name that is associated with the security key for the BFD session. The name you specify must match one of the key chains configured in the **authentication-key-chains** *key-chain* statement at the **[edit security]** hierarchy level.
- Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.

3. Configure BFD timers for aggregated Ethernet interfaces.

To specify the BFD timers, include the **detection-time** statement:

```
bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
}
```

Specify the threshold value. This is the maximum time interval for detecting a BFD neighbor. If the transmit interval is greater than this value, the device triggers a trap.

4. Configure 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.

To specify the hold-down interval, include the **holddown-interval** statement:

```
bfd-liveness-detection {
  holddown-interval milliseconds;
}
```

You can configure a number in the range from 0 through 255,000 milliseconds, and the default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

This value represents the minimum interval at which the local routing device transmits BFD packets, as well as the minimum interval in which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.

5. Configure the source address for the BFD session.

To specify a local address, include the **local-address** statement:

```
bfd-liveness-detection {
  local-address bfd-local-address;
}
```

The BFD local address is the loopback address of the source of the BFD session.

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

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.

6. Specify the minimum interval that indicates the time interval for transmitting and receiving data.

This value represents the minimum interval at which the local routing device transmits BFD packets, as well as the minimum interval in which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement:

```
bfd-liveness-detection {
  minimum-interval milliseconds;
}
```

NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

7. Specify only the minimum receive interval for failure detection by including the **minimum-receive-interval** statement:

```
bfd-liveness-detection {
  minimum-receive-interval milliseconds;
}
```

This value represents the minimum interval in which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds.

8. Specify the number of BFD packets that were not received by the neighbor that causes the originating interface to be declared down by including the **multiplier** statement:

```
bfd-liveness-detection {
  multiplier number;
}
```

The default value is 3. You can configure a number in the range from 1 through 255.

9. Configure the neighbor in a BFD session.

The neighbor address can be either an IPv4 or an IPv6 address.

To specify the next hop of the BFD session, include the **neighbor** statement:

```
bfd-liveness-detection {
  neighbor bfd-neighbor-address;
}
```

The BFD neighbor address is the loopback address of the remote destination of the BFD session.

NOTE: Beginning with Junos OS Release 16.1, you can also configure the AE interface address of the remote destination as the BFD neighbor address in a micro BFD session.

10. (Optional) Configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement:

```
bfd-liveness-detection {
  no-adaptation;
}
```

NOTE: We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

11. Specify a threshold for detecting the adaptation of the detection time by including the **threshold** statement:

```
bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
}
```

When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the minimum-interval or the minimum-receive-interval value. The threshold must be a higher value than the multiplier for either of these configured values. For example, if the minimum-receive-interval is 300 ms and the multiplier is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value greater than 900.

12. Specify only the minimum transmit interval for failure detection by including the **transmit-interval** **minimum-interval** statement:

```
bfd-liveness-detection {
  transmit-interval {
    minimum-interval milliseconds;
  }
}
```

This value represents the minimum interval at which the local routing device transmits BFD packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

13. Specify the transmit threshold for detecting the adaptation of the transmit interval by including the **transmit-interval threshold** statement:

```
bfd-liveness-detection {
  transmit-interval {
    threshold milliseconds;
  }
}
```

The threshold value must be greater than the transmit interval. When the BFD session detection time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the minimum-interval or the minimum-receive-interval value. The threshold must be a higher value than the multiplier for either of these configured values.

14. Specify the BFD version by including the **version** statement:

```
bfd-liveness-detection {
  version (1 | automatic);
}
```

The default is to have the version detected automatically.

NOTE:

- The **version** option is not supported on the QFX Series. Starting in Junos OS Release 17.2R1, a warning will appear if you attempt to use this command.
- This feature works when both the devices support BFD. If BFD is configured at only one end of the LAG, this feature does not work.

SEE ALSO

[authentication](#)[bfd-liveness-detection](#) | [561](#)[detection-time](#)*Example: Configuring Independent Micro BFD Sessions for LAG*

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

IN THIS SECTION

- [Understanding the Hashing Algorithm](#) | [343](#)
- [IP \(IPv4 and IPv6\)](#) | [344](#)
- [MPLS](#) | [346](#)
- [MAC-in-MAC Packet Hashing](#) | [347](#)
- [Layer 2 Header Hashing](#) | [348](#)
- [Hashing Parameters](#) | [349](#)

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

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

NOTE: Platform support depends on the Junos OS release in your installation.

This topic contains the following sections:

Understanding the Hashing Algorithm

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

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

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

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

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

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

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

Note the following points regarding the hashing algorithm:

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

MPLS

The hashing algorithm hashes MPLS packets using the source IP, destination IP, MPLS label 0, MPLS label 1, MPLS label 2, and MPLS 3 fields. On the QFX5110 and QFX5200 switches, LSR routers also support ECMP. ECMP uses these fields for hashing on an LSR router:

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

Table 56 on page 346 displays the MPLS payload fields that are used by the hashing algorithm, by default:

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

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

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

Table 56: MPLS Hashing Fields

Field	EX3400	EX4300	QFX5100	QFX5110	QFX5200
Source MAC	X	X	X	X	X
Destination MAC	X	X	X	X	X
EtherType	X	X	X	X	X
VLAN ID	X	X	X	X	X
Source IP	✓	✓	✓	✓	✓
Destination IP	✓	✓	✓	✓	✓
Protocol (for IPv4 packets)	X	X	X	X	X
Next header (for IPv6 packets)	X	X	X	X	X
Layer 4 Source Port	X	X	X	X	X

Table 56: MPLS Hashing Fields (*continued*)

Field	EX3400	EX4300		QFX5100	QFX5110	QFX5200
Layer 4 Destination Port	X	X		X	X	X
IPv6 Flow lab	X	X		X	X	X
MPLS label 0	X	✓		✓	✓	✓
MPLS label 1	✓	✓		✓	✓	✓
MPLS label 2	✓	✓		✓	✓	✓
MPLS label 3	✓	X		X	X	X
Ingress Port ID	✓ (LSR and L2Circuit)	X	X	X	✓ (LSR and L2Circuit)	✓ (LSR and L2Circuit)

MAC-in-MAC Packet Hashing

Packets using the MAC-in-MAC EtherType are hashed by the hashing algorithm using the Layer 2 payload source MAC, Layer 2 payload destination MAC, and Layer 2 payload EtherType fields. See [Table 57 on page 347](#).

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

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

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

Table 57: MAC-in-MAC Hashing Fields

Field	EX3400	EX4300	QFX5100	QFX5110	QFX5200
Layer 2 Payload Source MAC	✓	✓	✓	✓	✓

Table 57: MAC-in-MAC Hashing Fields (*continued*)

Field	EX3400	EX4300	QFX5100	QFX5110	QFX5200
Layer 2 Payload Destination MAC	✓	✓	✓	✓	✓
Layer 2 Payload EtherType	✓	✓	✓	✓	✓
Layer 2 Payload Outer VLAN	✓	X	X	X	X

Layer 2 Header Hashing

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

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

Table 58: Layer 2 Header Hashing Fields

Field	EX3400	EX4300	QFX5100	QFX5110	QFX5200
Source MAC	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
Destination MAC	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
EtherType	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
VLAN ID	X (configurable)	X (configurable)	X (configurable)	✓ (configurable)	✓ (configurable)

Hashing Parameters

Starting in Junos OS Release 19.1R1, on the QFX5000 line of switches, you can change hashing parameters for the existing algorithms implemented. You can change the threshold of shared buffer pools for both ingress and egress buffer partitions and you can make changes to the hash function selection, hash algorithm, and other additional parameters. See [Configuring the Shared-Buffer Threshold and Other Hashing Parameters](#), later in this document.

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

IN THIS SECTION

- [Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing | 350](#)
- [Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing | 350](#)
- [Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing | 351](#)
- [Configuring Other Hashing Parameters | 351](#)

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

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

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

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

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

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

1. Configure the hash mode to Layer 2 header:

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

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

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

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

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

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

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

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

Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing

To configure the hashing algorithm to use fields in the IP payload for hashing:

1. Configure the hash mode to Layer 2 payload:

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

The IP payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IP payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet {no-ipv4-destination-address | no-ipv4-source-address | no-l4-destination-port
| no-l4-source-port | no-protocol | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and protocol fields and instead hash traffic based only on the IPv4 source and destination addresses:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet no-l4-destination-port no-l4-source-port no-protocol
```

Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing

To configure the hashing algorithm to use fields in the IPv6 payload for hashing:

1. Configure the hash mode to Layer 2 payload:

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

The IPv6 payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IPv6 payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet6 {no-ipv6-destination-address | no-ipv6-source-address |
no-l4-destination-port | no-l4-source-port | no-next-header | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and the Next Header fields and instead hash traffic based only on the IPv6 source and IPv6 destination address fields only:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet6 no-l4-destination-port no-l4-source-port no-next-header
```

Configuring Other Hashing Parameters

To configure hashing parameters for either ECMP or LAG traffic:

1. Configure the preprocess parameter:

```
[edit forwarding-options enhanced-hash-key]
```

```
user@switch# set hash-parameters (ecmp | lag) preprocess
```

2. Configure the function parameter:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-parameters (ecmp | lag) function (crc16-bisync | crc16-ccitt | crc32-hi |
crc32-lo)
```

3. Configure the offset value:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-parameters (ecmp | lag) offset offset-value
```

SEE ALSO

[Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic | 342](#)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

Load Balancing for Aggregated Ethernet Interfaces

IN THIS SECTION

- [Load Balancing and Ethernet Link Aggregation Overview | 353](#)
- [Configuring Load Balancing Based on MAC Addresses | 353](#)
- [Configuring Load Balancing on a LAG Link | 355](#)
- [Example: Configuring Load Balancing on a LAG Link | 355](#)
- [Understanding Consistent Load Balancing Through Resilient Hashing on ECMP Groups | 356](#)
- [Configuring Consistent Load Balancing for ECMP Groups | 357](#)
- [Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches | 360](#)
- [Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches | 365](#)
- [Dynamic Load Balancing | 372](#)

- [Configuring Dynamic Load Balancing | 374](#)
- [Example: Configure Dynamic Load Balancing | 376](#)

Load balancing is done on Layer 2 across the member links making the configuration better without congestion and maintaining redundancy. The below topics discuss the overview of load balancing, configuring load balancing based on MAC addresses and on LAG link, understanding the consistency through resilient hashing.

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” on page 355](#). In a Layer 2 switch, one link is overutilized and other links are underutilized.

SEE ALSO

| *payload*

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include

the **multiservice** statement at the **[edit forwarding-options hash-key]** or **[edit chassis fpc slot number pic PIC number hash-key]** hierarchy level:

```
multiservice {  
  source-mac;  
  destination-mac;  
  payload {  
    ip {  
      layer3-only;  
      layer-3 (source-ip-only | destination-ip-only);  
      layer-4;  
      inner-vlan-id;  
      outer-vlan-id;  
    }  
  }  
}
```

To include the destination-address MAC information in the hash key, include the **destination-mac** option.
To include the source-address MAC information in the hash key, include the **source-mac** option.

NOTE: Any packets that have the same source and destination address will be sent over the same path.

NOTE: You can configure per-packet load balancing to optimize EVPN traffic flows across multiple paths.

NOTE: Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.

SEE ALSO

| [multiservice](#) | 626

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.

By default, Junos 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.

For more information about link aggregation group (LAG) configuration, see the *Junos OS Network Interfaces Library for Routing Devices*.

Example: Configuring Load Balancing on a LAG Link

This example configures the load-balancing hash key to use the source Layer 3 IP address option and Layer 4 header fields as well as the source and destination MAC addresses for load balancing on a link aggregation group (LAG) link:

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

NOTE: Any change in the hash key configuration requires a reboot of the FPC for the changes to take effect.

Understanding Consistent Load Balancing Through Resilient Hashing on ECMP Groups

You can use consistent load balancing to minimize flow remapping in an equal-cost multipath (ECMP) group.

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, Junos OS rechooses the next-hop address by using the hash algorithm.

You can configure *consistent load balancing* on the switch to prevent the reordering of *all* flows to active paths in an ECMP group when one or more next-hop paths fail. Only flows for paths that are inactive are redirected to another active next-hop path. Flows mapped to servers that remain active are maintained.

This feature applies only to external BGP peers.

Configuring Consistent Load Balancing for ECMP Groups

Per-packet load balancing allows you to spread traffic across multiple equal-cost paths. By default, when a failure occurs in one or more paths, the hashing algorithm recalculates the next hop for all paths, typically resulting in the redistribution of all flows. *Consistent load balancing* enables you to override this behavior so that only flows for links that are inactive are redirected. All existing active flows are maintained without disruption. In a data center environment, the redistribution of all flows when a link fails potentially results in significant traffic loss or a loss of service to servers whose links remain active. Consistent load balancing maintains all active links and instead remaps only those flows affected by one or more link failures. This feature ensures that flows connected to links that remain active continue uninterrupted.

This feature applies to topologies where members of an equal-cost multipath (ECMP) group are external BGP neighbors in a single-hop BGP session. Consistent load balancing does not apply when you add a new ECMP path or modify an existing path in any way. To add a new path with minimal disruption, define a new ECMP group without modifying the existing paths. In this way, clients can be moved to the new group gradually without terminating existing connections.

- (On MX Series) Only Modular Port Concentrators (MPCs) are supported.
- Both IPv4 and IPv6 paths are supported.
- ECMP groups that are part of a virtual routing and forwarding (VRF) instance or other routing instance are also supported.
- Multicast traffic is not supported.
- Aggregated interfaces are supported, but consistent load balancing is not supported among members of the link aggregation (LAG) bundle. Traffic from active members of the LAG bundle might be moved to another active member when one or more member links fail. Flows are rehashed when one or more LAG member links fail.
- We strongly recommend that you apply consistent load balancing to no more than a maximum of 1,000 IP prefixes per router or switch.
- Layer 3 adjacency over integrated routing and bridging (IRB) interfaces is supported.

You can configure the BGP *add-path* feature to enable replacement of a failed path with a new active path when one or more paths in the ECMP group fail. Configuring replacement of failed paths ensures that traffic flow on the failed paths only are redirected. Traffic flow on active paths will remain unaltered.

NOTE:

- When you configure consistent load balancing on generic routing encapsulation (GRE) tunnel interfaces, you must specify the inet address of the far end GRE interface so that the Layer 3 adjacencies over the GRE tunnel interfaces are installed correctly in the forwarding table. However, ECMP fast reroute (FRR) over GRE tunnel interfaces is not supported during consistent load balancing. You can specify the destination address on the router configured with consistent load balancing at the **[edit interfaces interface name unit unit name family inet address address]** hierarchy level. For example:

```
[edit interfaces]
user@host# set interfaces gr-4/0/0 unit 21 family inet address 10.10.31.2/32 destination 10.10.31.1
```

For more information on generic routing encapsulation see [“Configuring Generic Routing Encapsulation Tunneling” on page 255](#).

- Consistent load balancing does not support BGP multihop for EBGP neighbors. Therefore, do not enable the **multihop** option on devices configured with consistent load balancing.

To configure consistent load balancing for ECMP groups:

1. Configure BGP and enable the BGP group of external peers to use multiple paths.
2. Create a routing policy to match incoming routes to one or more destination prefixes.

```
[edit policy-options]
user@host# set policy-statement policy-statement-name from route-filter destination-prefix orlonger
```

3. Apply consistent load balancing to the routing policy so that only traffic flows to one or more destination prefixes that experience a link failure are redirected to an active link.

```
[edit policy-options]
user@host# set policy-statement policy-statement-name then load-balance consistent-hash
```

4. Create a separate routing policy and enable per-packet load balancing.

NOTE: You must configure and apply a per-packet load-balancing policy to install all routes in the forwarding table.

```
[edit policy-options]
user@host# set policy-statement policy-statement-name then load-balance per-packet
```

5. Apply the routing policy for consistent load balancing to the BGP group of external peers.

NOTE: Consistent load balancing can be applied only to BGP external peers. This policy cannot be applied globally.

```
[edit protocols bgp]
user@host# set group group-name import policy-statement-name
#This policy-statement-name refers to the policy created in Step 2.
```

6. (Optional) Enable bidirectional forwarding detection (BFD) for each external BGP neighbor.

```
[edit protocols bgp]
user@host# set group group-name neighbor ip-address bfd-liveness-detection milliseconds
```

NOTE: This step shows the minimum BFD configuration required. You can configure additional options for BFD.

7. Apply the per-prefix load-balancing policy globally to install all next-hop routes in the forwarding table.

```
[edit routing-options]
user@host# set forwarding-table export policy-statement-name
#This policy-statement-name refers to the policy created in Step 4.
```

8. (Optional) Enable fast reroute for ECMP routes.

```
[edit routing-options]
user@host# set forwarding-table ecmp-fast-reroute
```

9. Verify the status of one or more ECMP routes for which you enabled consistent load balancing.

```
user@host> show route destination-prefix extensive
```

The output of the command displays the following flag when consistent load balancing is enabled:
State: <Active Ext LoadBalConsistentHash>

SEE ALSO

[policy-statement | 708](#)

Actions in Routing Policy Terms

[Understanding Per-Packet Load Balancing | 257](#)

Examples: Configuring BGP Multipath

Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches

IN THIS SECTION

- [Create LAGs for Multicasting in Increments of 10 Gigabits | 361](#)
- [When Should I Use Multicast Load Balancing? | 362](#)
- [How Does Multicast Load Balancing Work? | 363](#)
- [How Do I Implement Multicast Load Balancing on an EX8200 Switch? | 364](#)

Streaming video technology was introduced in 1997. Multicast protocols were subsequently developed to reduce data replication and network overloads. With multicasting, servers can send a single stream to a group of recipients instead of sending multiple unicast streams. While the use of streaming video technology was previously limited to occasional company presentations, multicasting has provided a boost to the technology resulting in a constant stream of movies, real-time data, news clips, and amateur videos flowing nonstop to computers, TVs, tablets, and phones. However, all of these streams quickly overwhelmed the capacity of network hardware and increased bandwidth demands leading to unacceptable blips and stutters in transmission.

To satisfy the growing bandwidth demands, multiple links were virtually aggregated to form bigger logical point-to-point link channels for the flow of data. These virtual link combinations are called multicast interfaces, also known as link aggregation groups (LAGs).


Multicast load balancing involves managing the individual links in each LAG to ensure that each link is used efficiently. Hashing algorithms continually evaluate the data stream, adjusting stream distribution over the links in the LAG, so that no link is underutilized or overutilized. Multicast load balancing is enabled by default on Juniper Networks EX8200 Ethernet Switches.

This topic includes:

Create LAGs for Multicasting in Increments of 10 Gigabits

The maximum link size on an EX8200 switch is 10 gigabits. If you need a larger link on an EX8200 switch, you can combine up to twelve 10-gigabit links. In the sample topology shown in [Figure 10 on page 362](#), four 10-gigabit links have been aggregated to form each 40-gigabit link.

Figure 10: 40-Gigabit LAGs on EX8200 Switches



When Should I Use Multicast Load Balancing?

Use a LAG with multicast load balancing when you need a downstream link greater than 10 gigabits. This need frequently arises when you act as a service provider or when you multicast video to a large audience.

To use multicast load balancing, you need the following:

- An EX8200 switch—Standalone switches support multicast load balancing, while Virtual Chassis does not.
- A Layer 3 routed multicast setup—For information about configuring multicasting, see [Junos OS Routing Protocols Configuration Guide](#).

- Aggregated 10-gigabit links in a LAG—For information about configuring LAGs with multicast load balancing, see *Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)*.

How Does Multicast Load Balancing Work?

Juniper Networks Junos operating system (Junos OS) supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs and is supported only on Layer 3 interfaces. When traffic can use multiple member links, traffic that is part of the same stream must always be on the same link.

Multicast load balancing uses one of seven available hashing algorithms and a technique called queue shuffling (alternating between two queues) to distribute and balance the data, directing streams over all available aggregated links. You can select one of the seven algorithms when you configure multicast load balancing, or you can use the default algorithm, `crc-sgip`, which uses a cyclic redundancy check (CRC) algorithm on the multicast packets' group IP address. We recommend that you start with the `crc-sgip` default and try other options if this algorithm does not evenly distribute the Layer 3 routed multicast traffic. Six of the algorithms are based on the hashed value of IP addresses (IPv4 or IPv6) and will produce the same result each time they are used. Only the balanced mode option produces results that vary depending on the order in which streams are added. See [Table 59 on page 363](#) for more information.

Table 59: Hashing Algorithms Used by Multicast Load Balancing

Hashing Algorithms	Based On	Best Use
<code>crc-sgip</code>	Cyclic redundancy check of multicast packets' source and group IP address	Default—high-performance management of IP traffic on 10-Gigabit Ethernet network. Predictable assignment to the same link each time. This mode is complex but yields a good distributed hash.
<code>crc-gip</code>	Cyclic redundancy check of multicast packets' group IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>crc-sip</code>	Cyclic redundancy check of multicast packets' source IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the stream sources vary.
<code>simple-sgip</code>	XOR calculation of multicast packets' source and group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sgip</code> yields. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic.

Table 59: Hashing Algorithms Used by Multicast Load Balancing (*continued*)

Hashing Algorithms	Based On	Best Use
simple-gip	XOR calculation of multicast packets' group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as crc-gip yields. Try this when crc-gip does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
simple-sip	XOR calculation of multicast packets' source IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as crc-sip yields. Try this mode when crc-sip does not evenly distribute the Layer 3 routed multicast traffic and stream sources vary.
balanced	Round-robin calculation method used to identify multicast links with the least amount of traffic	Best balance is achieved, but you cannot predict which link will be consistently used because that depends on the order in which streams come online. Use when consistent assignment is not needed after every reboot.

How Do I Implement Multicast Load Balancing on an EX8200 Switch?

To implement multicast load balancing with an optimized level of throughput on an EX8200 switch, follow these recommendations:

- Allow 25 percent unused bandwidth in the aggregated link to accommodate any dynamic imbalances due to link changes caused by sharing multicast interfaces.
- For downstream links, use multicast interfaces of the same size whenever possible. Also, for downstream aggregated links, throughput is optimized when members of the aggregated link belong to the same devices.
- For upstream aggregated links, use a Layer 3 link whenever possible. Also, for upstream aggregated links, throughput is optimized when the members of the aggregated link belong to different devices.

SEE ALSO

Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)

Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches

IN THIS SECTION

- [Requirements | 365](#)
- [Overview and Topology | 366](#)
- [Configuration | 367](#)
- [Verification | 370](#)

EX8200 switches support multicast load balancing on link aggregation groups (LAGs). Multicast load balancing evenly distributes Layer 3 routed multicast traffic over the LAGs. You can aggregate up to twelve 10-gigabit Ethernet links to form a 120-gigabit virtual link or LAG. The MAC client can treat this virtual link as if it were a single link to increase bandwidth, provide graceful degradation as link failures occur, and increase availability. On EX8200 switches, multicast load balancing is enabled by default. However, if it is explicitly disabled, you can reenabling it.

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

NOTE: Only EX8200 standalone switches with 10-gigabit links support multicast load balancing. Virtual Chassis does not support multicast load balancing.

This example shows how to configure a LAG and reenabling multicast load balancing:

Requirements

This example uses the following hardware and software components:

- Two EX8200 switches, one used as the access switch and one used as the distribution switch
- Junos OS Release 12.2 or later for EX Series switches

Before you begin:

- Configure four 10-gigabit interfaces on the EX8200 distribution switch: xe-0/1/0, xe-1/1/0, xe-2/1/0, and xe-3/1/0. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*.

Overview and Topology


Multicast load balancing uses one of seven hashing algorithms to balance traffic between the individual 10-gigabit links in the LAG. For a description of the hashing algorithms, see [multicast-loadbalance](#). The default hashing algorithm is crc-sgip. You can experiment with the different hashing algorithms until you determine the one that best balances your Layer 3 routed multicast traffic.

When a link larger than 10 gigabits is needed on an EX8200 switch, you can combine up to twelve 10-gigabit links to create more bandwidth. This example uses the link aggregation feature to combine four 10-gigabit links into a 40-gigabit link on the distribution switch. In addition, multicast load balancing is enabled to ensure even distribution of Layer 3 routed multicast traffic on the 40-gigabit link. In the sample topology illustrated in [Figure 11 on page 367](#), an EX8200 switch in the distribution layer is connected to an EX8200 switch in the access layer.

NOTE: Link speed is automatically determined based on the size of the LAG configured. For example, if a LAG is composed of four 10-gigabit links, the link speed is 40 gigabits per second).

NOTE: The default hashing algorithm, crc-sgip, involves a cyclic redundancy check of both the multicast packet source and group IP addresses.

Figure 11: 40-Gigabit LAG Composed of Four 10-Gigabit Links



You will configure a LAG on each switch and reenabling multicast load balancing. When reenabled, multicast load balancing will automatically take effect on the LAG, and the speed is set to 10 gigabits per second for each link in the LAG. Link speed for the 40-gigabit LAG is automatically set to 40 gigabits per second.

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.

```
set chassis aggregated-devices ethernet device-count 1
```

```
set interfaces ae0 aggregated-ether-options minimum-links 1
```

```
set interfaces xe-0/1/0 ether-options 802.3ad ae0
```

```
set interfaces xe-1/1/0 ether-options 802.3ad ae0
```

```
set interfaces xe-2/1/0 ether-options 802.3ad ae0
```

```
set interfaces xe-3/1/0 ether-options 802.3ad ae0
```

```
set chassis multicast-loadbalance hash-mode crc-gip
```

Step-by-Step Procedure

To configure a LAG and reenenable multicast load balancing:

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

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

2. Specify the minimum number of links for the aggregated Ethernet interface (aex), that is, the LAG, to be labeled **up**:

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

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links 1
```

3. Specify the four members to be included within the LAG:

```
[edit interfaces]
user@switch# set xe-0/1/0 ether-options 802.3ad ae0
user@switch# set xe-1/1/0 ether-options 802.3ad ae0
user@switch# set xe-2/1/0 ether-options 802.3ad ae0
user@switch# set xe-3/1/0 ether-options 802.3ad ae0
```

4. Reenable multicast load balancing:

```
[edit chassis]
user@switch# set multicast-loadbalance
```


NOTE: You do not need to set link speed the way you do for LAGs that do not use multicast load balancing. Link speed is automatically set to 40 gigabits per second on a 40-gigabit LAG.

5. You can optionally change the value of the **hash-mode** option in the **multicast-loadbalance** statement to try different algorithms until you find the one that best distributes your Layer 3 routed multicast traffic.

If you change the hashing algorithm when multicast load balancing is disabled, the new algorithm takes effect after you reenables multicast load balancing.

Results

Check the results of the configuration:

```
user@switch> show configuration
chassis
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
  multicast-loadbalance {
    hash-mode crc-gip;
  }

interfaces
  xe-0/1/0 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-1/1/0 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-2/1/0 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-3/1/0 {
```

```
ether-options {
  802.3ad ae0;
}
ae0 {
  aggregated-ether-options {
    minimum-links 1;
  }
}
```

Verification

IN THIS SECTION

- [Verifying the Status of a LAG Interface | 370](#)
- [Verifying Multicast Load Balancing | 371](#)

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

Verifying the Status of a LAG Interface

Purpose

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

Action

Verify that the **ae0** LAG has been created:

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

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

Meaning

The interface name aex indicates that this is a LAG. A stands for aggregated, and E stands for Ethernet. The number differentiates the various LAGs.

Verifying Multicast Load Balancing

Purpose

Check that traffic is load-balanced equally across paths.

Action

Verify load balancing across the four interfaces:

```
user@switch> monitor interface traffic
```

```
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
ibmoem02-rel                      Seconds: 3                      Time: 16:06:14
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
xe-0/1/0	Up	2058834	(10)	7345862	(19)
xe-1/1/0	Up	2509289	(9)	6740592	(21)
xe-2/1/0	Up	8625688	(90)	10558315	(20)
xe-3/1/0	Up	2374154	(23)	71494375	(9)

Meaning

The interfaces should be carrying approximately the same amount of traffic.

SEE ALSO

Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)

Dynamic Load Balancing

Load balancing is used to ensure that network traffic is distributed as evenly as possible across members in a given ECMP (Equal-cost multi-path routing) or LAG (Link Aggregation Group). In general, load balancing is classified as either static or dynamic. Static load balancing (SLB) computes hashing solely based on the packet contents (for example, source IP, destination IP, and so on.). The biggest advantage of SLB is that packet ordering is guaranteed as all packets of a given flow take the same path. However, because the SLB mechanism does not consider the path or link load, the network often experiences the following problems:

- Poor link bandwidth utilization
- Elephant flow on a single link completely dropping mice flows on it.

Dynamic load balancing (DLB) is an improvement on top of SLB.

For ECMP, you can configure DLB globally, whereas for LAG, you configure it for each aggregated Ethernet interface. You can apply DLB on selected **ether-type** (IPv4, IPv6, and MPLS) based on configuration. If you don't configure any **ether-type**, then DLB is applied to all EtherTypes. Note that you must explicitly configure the DLB mode because there is no default mode.

NOTE:

- Starting in Junos OS Release 19.4R1, QFX5120-32C and QFX5120-48Y switches support dynamic load balancing for both ECMP and LAG. For LAG, DLB must be configured on per aggregated ethernet interface basis.
- Starting in Junos OS evolved Release 19.4R2, QFX5220 switches support dynamic load balancing (DLB) for ECMP. For ECMP, DLB must be configured globally.
- You cannot configure both DLB and resilient hashing at the same time. Otherwise, a commit error will be thrown.
- DLB is applicable only for unicast traffic.
- DLB is not supported when the LAG is one of the egress ECMP members.
- DLB is not supported for remote LAG members.
- DLB is not supported on Virtual Chassis and Virtual Chassis Fabric (VCF).
- DLB on LAG and HiGig-trunk are not supported at the same time.
- QFX5220 switches does not support DLB on LAG.

Table 60: Platforms That Support Dynamic Load Balancing for ECMP/LAG

Platform	DLB Support for ECMP	DLB Support for LAG
----------	----------------------	---------------------

Table 60: Platforms That Support Dynamic Load Balancing for ECMP/LAG (*continued*)

Platform	DLB Support for ECMP	DLB Support for LAG
QFX5120-32C	Yes	Yes
QFX5120-48Y	Yes	Yes
QFX5220	Yes	No

You can use the following DLB modes to load-balance traffic:

- *Per packet mode*

In this mode, DLB is initiated for each packet in the flow. This mode makes sure that the packet always gets assigned to the best-quality member port. However, in this mode, DLB may experience packet reordering problems that can arise due to latency skews.

- *Flowlet mode*

This mode relies on assigning links based on *flowlets* instead of flows. Real-world application traffic relies on flow control mechanisms of upper-layer transport protocols such as TCP, which throttle the transmission rate. As a result, flowlets are created. You can consider flowlets as multiple bursts of the same flow separated by a period of inactivity between these bursts—this period of inactivity is referred to as the inactivity interval. The inactivity interval serves as the demarcation criteria for identifying new flowlets and is offered as a user-configurable statement under the DLB configuration. In this mode, DLB is initiated per flowlet—that is, for the new flow as well as for the existing flow that has been inactive for a sufficiently long period of time (configured **inactivity-interval**). The reordering problem of per packet mode is addressed in this mode as all the packets in a flowlet take the same link. If the **inactivity-interval** value is configured to be higher than the maximum latency skew across all ECMP paths, then you can avoid packet reordering across flowlets while increasing link utilization of all available ECMP links.

- *Assigned flow mode*

You can use assigned flow mode to selectively disable rebalancing for a period of time to isolate problem sources. You cannot use this mode for real-time DLB or predict the egress ports that will be selected using this mode because assigned flow mode does not consider port load and queue size.

NOTE: Here are some of the important behaviors of DLB:

- DLB is applicable for incoming EtherTypes only.
- From a DLB perspective, both Layer 2 and Layer 3 link aggregation group (LAG) bundles are considered the same.
- The link utilisation will not be optimal if you use dynamic load balancing in asymmetric bundles—that is, on ECMP links with different member capacities.
- With DLB, no reassignment of flow happens when a new link is added in per packet and assigned flow modes. This can cause suboptimal usage in link flap scenarios where a utilized link may not be utilized after it undergoes a flap if no new flow or flowlets are seen after the flap.

Benefits

- DLB considers member bandwidth utilization along with packet content for member selection. As a result, we achieve better link utilization based on real-time link loads.
- DLB ensures that links hogged by elephant flows are not used by mice flows. Thus, by using DLB, we avoid hash collision drops that occur with SLB. That is, with DLB the links are spread across, and thus the collision and the consequent drop of packets are avoided.

SEE ALSO

[Configuring Dynamic Load Balancing | 374](#)

[Example: Configure Dynamic Load Balancing | 376](#)

[dlb | 575](#)

[show forwarding-options enhanced-hash-key | 818](#)

Configuring Dynamic Load Balancing

This topic describes how to configure dynamic load balancing (DLB) in flowlet mode.

Starting in Junos OS Release 19.4R1, QFX5120-32C and QFX5120-48Y switches support dynamic load balancing for both ECMP and LAG. For LAG, DLB must be configured on per aggregated ethernet interface basis.

Starting in Junos OS evolved Release 19.4R2, QFX5220 switches support dynamic load balancing (DLB) for ECMP. For ECMP, DLB must be configured globally.

Configuring DLB for ECMP (Flowlet mode)

To configure dynamic load balancing for ECMP with flowlet mode (QFX5120-32C, QFX5120-48Y, and QFX5220 switches):

1. Enable dynamic load balancing with flowlet mode:

```
[edit forwarding-options enhanced-hash-key]
user@router# set ecmp-dlb flowlet
```

2. (Optional) Configure the *inactivity-interval* value - minimum inactivity interval (in micro seconds) for link re-assignment:

```
[edit forwarding-options enhanced-hash-key]
user@router# set ecmp-dlb flowlet inactivity-interval (micro seconds)
```

3. (Optional) Configure dynamic load balancing with **ether-type**:

```
[edit forwarding-options enhanced-hash-key]
user@router# set ecmp-dlb ether-type mpls
```

4. (Optional) You can view the options configured for dynamic load balancing on ECMP using **show forwarding-options enhanced-hash-key** command.

Similarly, you can configure DLB for ECMP with *Per packet* or *Assigned flow* mode.

Configuring DLB for LAG (Flowlet mode)

Before you begin, create an aggregated ethernet (AE) bundle by configuring a set of router interfaces as aggregated Ethernet and with a specific aggregated ethernet (AE) group identifier.

To configure dynamic load balancing for LAG with flowlet mode (QFX5120-32C and QFX5120-48Y):

1. Enable dynamic load balancing with flowlet mode:

```
[edit interfaces ae-x aggregated-ether-options]
user@router# set dlb flowlet
```

2. (Optional) Configure the *inactivity-interval* value - minimum inactivity interval (in micro seconds) for link re-assignment:

```
[edit interfaces ae-x aggregated-ether-options]
user@router# set dlb flowlet inactivity-interval (micro seconds)
```

3. (Optional) Configure dynamic load balancing with **ether-type**:

```
[edit forwarding-options enhanced-hash-key]
user@router# set lag-dlb ether-type mpls
```

4. (Optional) You can view the options configured for dynamic load balancing on LAG using **show forwarding-options enhanced-hash-key** command.

Similarly, you can configure DLB for LAG with *Per packet* or *Assigned flow* mode.

SEE ALSO

[Dynamic Load Balancing | 372](#)

[Example: Configure Dynamic Load Balancing | 376](#)

[dlb | 575](#)

[show forwarding-options enhanced-hash-key | 818](#)

Example: Configure Dynamic Load Balancing

IN THIS SECTION

- [Requirements | 377](#)
- [Overview | 377](#)
- [Configuration | 377](#)
- [Verification | 381](#)

This example shows how to configure dynamic load balancing.

Requirements

This example uses the following hardware and software components:

- Two QFX5120-32C or QFX5120-48Y switches
- Junos OS Release 19.4R1 or later running on all devices

Overview

Dynamic load balancing (DLB) is an improvement on top of SLB.


For ECMP, you can configure DLB globally, whereas for LAG, you configure it for each aggregated Ethernet interface. You can apply DLB on selected **ether-type** such as IPv4, IPv6, and MPLS based on configuration. If you don't configure any **ether-type**, then DLB is applied to all EtherTypes. Note that you must explicitly configure the DLB mode because there is no default mode.

NOTE:

- Starting in Junos OS Release 19.4R1, QFX5120-32C and QFX5120-48Y switches support dynamic load balancing on both ECMP and LAG.
- You cannot configure both DLB and Resilient Hashing at the same time. Otherwise, commit error will be thrown.

Topology

In this topology, both R0 and R1 are connected.



NOTE: This example shows static configuration. You can also add configuration with dynamic protocols.

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.

R0

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.2/24
set interfaces xe-0/0/10 unit 0 family inet address 10.1.1.2/24
set interfaces xe-0/0/54:0 unit 0 family inet address 10.10.10.2/24
set forwarding-options enhanced-hash-key ecmp-dlb per-packet
set policy-options policy-statement loadbal then load-balance per-packet
set routing-options static route 20.0.1.0/24 next-hop 10.1.0.3
set routing-options static route 20.0.1.0/24 next-hop 10.1.1.3
set routing-options forwarding-table export loadbal
```

R1

```
set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.3/24
set interfaces xe-0/0/10 unit 0 family inet address 10.1.1.3/24
set interfaces xe-0/0/52:0 unit 0 family inet address 20.0.0.2/16
```

Configure Dynamic Load Balancing for LAG (QFX5120-32C and QFX5120-48Y)

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 R0 router:

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

1. Configure Link Aggregation Group (LAG).

```
[edit interfaces]
user@R0# set interfaces xe-0/0/0 ether-options 802.3ad ae0
user@R0# set interfaces xe-0/0/10 ether-options 802.3ad ae0
```

```

user@R0# set interfaces ae0 aggregated-ether-options lacp active
user@R0# set interfaces ae0 unit 0 family inet address 10.1.0.2/24
user@R0# set routing-options static route 20.0.1.0/24 next-hop 10.1.0.3

```

After configuring LAG, in the verification section, execute the steps in the *Verifying Traffic Load before configuring Dynamic Load Balancing Feature on LAG* section, to check the configuration or the traffic load before configuring DLB.

2. Configure Dynamic Load Balancing with per-packet mode for LAG.

```

[edit]
user@R0# set interfaces ae0 aggregated-ether-options dlb per-packet

```

After configuring the DLB, in the verification section, execute the steps in the *Verifying Traffic Load after configuring Dynamic Load Balancing Feature on LAG* section, to check the configuration or the traffic load before configuring DLB.

Verification

IN THIS SECTION

- [Verify Traffic Load Before Configuring Dynamic Load Balancing Feature on LAG | 379](#)
- [Verify Traffic Load After Configuring Dynamic Load Balancing Feature on LAG | 380](#)

Confirm that the configuration is working properly.

Verify Traffic Load Before Configuring Dynamic Load Balancing Feature on LAG

Purpose

Verify before the DLB feature is configured on the Link Aggregation Group.

Action

From operational mode, run the **show interfaces interface-name | match pps** command.

```
user@R0>show interfaces xe-0/0/0 | match pps
```

```

Input rate      : 1240 bps (1 pps)
Output rate     : 1024616 bps (1000 pps) ## all traffic in one link.

```

```
user@R0>show interfaces xe-0/0/10 | match pps
```

```

Input rate      : 616 bps (0 pps)
Output rate     : 1240 bps (1 pps)<<  Output rate      : 1240 bps (1 pps) ## no
traffic

```

Verify Traffic Load After Configuring Dynamic Load Balancing Feature on LAG

Purpose

Verify that packets received on the R0 are load-balanced.

Action

From operational mode, run the **show interfaces *interface-name*** command.

```
user@R0>show interfaces xe-0/0/0 | match pps
```

```

Input rate      : 616 bps (0 pps)
Output rate     : 519096 bps (506 pps)<<  Output rate      : 519096 bps (506 pps)
## load equally shared

```

```
user@R0>show interfaces xe-0/0/10 | match pps
```

```

Input rate      : 1232 bps (1 pps)
Output rate     : 512616 bps (500 pps)<<  Output rate      : 512616 bps (500 pps)
## load equally shared

```

Meaning

Dynamic Load balancing with per-packet mode successfully working. After applying dynamic load balancing feature on LAG, the load is equally shared in the network.

Configure Dynamic Load Balancing for ECMP (QFX5120-32C, QFX5120-48Y, and QFX5220 switches)

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 R0 router:

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

1. Configure the Gigabit Ethernet interface link connecting from R0 to R1.

[edit interfaces]

```

user@R0# set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.2/24
user@R0# set interfaces xe-0/0/10 unit 0 family inet address 10.1.1.2/24
user@R0# set interfaces xe-0/0/54:0 unit 0 family inet address 10.10.10.2/24

```

2. Create the static routes:

[edit interfaces]

```

user@R0# set routing-options static route 20.0.1.0/24 next-hop 10.1.0.3
user@R0# set routing-options static route 20.0.1.0/24 next-hop 10.1.1.3

```

3. Apply the load-balancing policy. The dynamic load balancing feature requires the multiple ECMP next hops to be present in the forwarding table.

[edit interfaces]

```

user@R0# set policy-options policy-statement loadbal then load-balance per-packet
user@R0# set routing-options forwarding-table export loadbal

```

4. Configure Dynamic Load Balancing with per-packet mode for ECMP.

[edit interfaces]

```

user@R0# set forwarding-options enhanced-hash-key ecmp-dlb per-packet

```

5. On R1, configure the Gigabit Ethernet interface link.

[edit interfaces]

```

user@R2# set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.3/24
user@R2# set interfaces xe-0/0/10 unit 0 family inet address 10.1.1.3/24
user@R2# set interfaces xe-0/0/52:0 unit 0 family inet address 20.0.0.2/16

```

Verification

IN THIS SECTION

- [Verify Dynamic Load Balancing on R0 | 382](#)
- [Verify Load Balancing on R1 | 382](#)

Confirm that the configuration is working properly at R0.

Verify Dynamic Load Balancing on R0

Purpose

Verify that packets received on the R0 are load-balanced.

Action

From operational mode, run the **run show route forwarding-table destination *destination-address*** command.

```
user@R0>show route forwarding-table destination 20.0.1.0/24
```

```
inet.0: 178 destinations, 178 routes (178 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.0.1.0/24          *[Static/5] 1d 03:35:12
                    > to 10.1.0.3 via xe-0/0/0.0
                    to 10.1.1.3 via xe-0/0/10.0
```

```
user@R0>show route 20.0.1.0/24
```

```
inet.0: 178 destinations, 178 routes (178 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.0.1.0/24          *[Static/5] 1d 03:35:12
                    > to 10.1.0.3 via xe-0/0/0.0
                    to 10.1.1.3 via xe-0/0/10.0
```

Meaning

Verify Load Balancing on R1

Purpose

Confirm that the configuration is working properly at R1.

Action

From operational mode, run the **show route** command.

```
user@R1>show route 20.0.1.25
```

```
inet.0: 146 destinations, 146 routes (146 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
20.0.0.0/16      *[Direct/0] ld 03:37:11  
                  >  via xe-0/0/52:0.0
```

Meaning

Dynamic Load balancing with per-packet mode successfully working. After applying dynamic load balancing feature on ECMP, the load is equally shared in the network.

SEE ALSO

[Dynamic Load Balancing | 372](#)

[Configuring Dynamic Load Balancing | 374](#)

[dlb | 575](#)

[show forwarding-options enhanced-hash-key | 818](#)

3

CHAPTER

Flexible Ethernet Services Encapsulation

Flexible Ethernet Services Encapsulation | 385

Flexible Ethernet Services Encapsulation

IN THIS SECTION

- [Understanding Flexible Ethernet Services Encapsulation on Switches | 385](#)
- [Configuring Flexible Ethernet Services Encapsulation to Support the Service Provider and Enterprise Styles of Configuration | 388](#)

Flexible Ethernet services is a type of encapsulation that enables a physical interface to support different types of Ethernet encapsulations at the logical interface level. Flexible Ethernet services encapsulation can be configured to support the service provided and the enterprise style configuration. The below topics discuss the overview of flexible Ethernet services encapsulation, and its configuration details.

Understanding Flexible Ethernet Services Encapsulation on Switches

IN THIS SECTION

- [Service Provider Style | 386](#)
- [Enterprise Style | 386](#)
- [Flexible Ethernet Services | 387](#)

Junos OS supports two different styles of configuration for switch interfaces: the service provider style and the enterprise style. The service provider style requires more configuration but provides greater flexibility. The enterprise style is easier to configure but offers less functionality. Each configuration style requires a different Ethernet encapsulation type. You can configure a physical interface to support both styles of configuration using flexible Ethernet services.

NOTE: On EX4300, QFX5100 (running Junos OS 16.1R5 or earlier), and QFX5200, the service provider style and enterprise style interface configurations are handled differently within Junos OS. If the service provider style and enterprise style interface configurations are mixed, the egress VLAN translation within the hardware can be incorrectly programmed leading to forwarding issues across the configured ports. Use the service provider style configuration in a Q-in-Q scenario. For all other scenarios, use the enterprise style configuration.

Flexible Ethernet services is a type of encapsulation that enables a physical interface to support different types of Ethernet encapsulations at the logical interface level. Defining multiple per-unit Ethernet encapsulations makes it easier to customize Ethernet-based services to multiple hosts connected to the same physical interface.

Service Provider Style

The service provider style of configuration allows for customization of Ethernet-based services at the logical interface level. Service providers typically have multiple customers connected to the same physical interface. Using the service provider style, you can configure multiple logical interfaces on the physical interface, and associate each unit with a different VLAN. This provides the flexibility to configure different services for each customer, but also requires more configuration, because each feature must be explicitly configured on the logical interface.

When configuring a physical interface to support only the service provider style, the physical interface must be encapsulated with the **extended-vlan-bridge** option to support bridging features. VLAN tagging must also be configured on the physical interface so that it can operate in trunk mode and transmit Ethernet frames with VLAN tags for multiple VLANs. Each logical interface is bound to a unique VLAN ID.

Enterprise Style

The enterprise style of configuration is designed to provide basic bridging functionality for consumers of Ethernet-based services. The isolation of services for different customers on a single port is not required, because each port is typically connected to a host or is providing a trunk to another switch.

With the enterprise style of configuration, logical interfaces are placed into Layer 2 mode by specifying **ethernet-switching** as the interface family. Without using flexible Ethernet services, **ethernet-switching** can only be configured on a single logical unit, unit 0. You cannot bind a VLAN ID to unit 0, because these interfaces operate either in trunk mode, which supports traffic with various VLAN tags, or in access mode, which supports untagged traffic.

Flexible Ethernet Services

The flexible Ethernet services encapsulation type enables a physical interface to support both styles of configuration. To support the service provider style, flexible Ethernet services allows for encapsulations to be configured at the logical interface level instead of the physical interface. To support the enterprise style, flexible Ethernet services allows the **ethernet-switching** family to be configured on any logical interface unit number instead of only unit 0.

For example, the configuration below shows three logical interfaces configured on a physical interface, xe-0/0/51, that is encapsulated for flexible Ethernet services. Unit 100 and unit 200 are configured in the service provider style and unit 300 is configured in the enterprise style. The encapsulation type of **vlan-bridge** is used to enable bridging on unit 100 and unit 200, and **family ethernet-switching** enables bridging on unit 300.

```

interfaces {
  xe-0/0/51 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 100 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
    unit 200 {
      encapsulation vlan-bridge;
      vlan-id 200;
    }
    unit 300 {
      family ethernet-switching {
        interface-mode trunk;
        vlan members 300;
      }
    }
  }
}

```

Following are the guidelines to follow when you configure the flexible Ethernet services encapsulation:

- On the QFX10000 line of switches, configuring either **vlan-tagging** or **flexible-vlan-tagging** with **family ethernet-switching** on the same interface is not supported.
- Only on the QFX10000 and EX9200 line of switches, you can enable **vlan-ccc** encapsulation when **flexible-ethernet-services** is already enabled.

- On QFX5100 switches, you can combine encapsulations on the same physical interface for **vlan-bridge** and **family ethernet switching**. Starting with Junos OS Release 16.1R6, you can also combine encapsulations on the same physical interface for **family inet** and **family ethernet-switching**.
- It is not required that the unit number and VLAN ID match, but it is considered a best practice.

Configuring Flexible Ethernet Services Encapsulation to Support the Service Provider and Enterprise Styles of Configuration

Flexible Ethernet services is a type of encapsulation that enables a physical interface to specify Ethernet encapsulations at the logical interface level. Each logical interface can have a different Ethernet encapsulation. Defining multiple per-unit Ethernet encapsulations makes it easier to customize Ethernet-based services to multiple hosts connected to the same physical interface.

An Ethernet interface that is not encapsulated with flexible Ethernet services and is operating in Layer 2 mode is limited to a single logical interface unit (0). Bridging is enabled on the interface by configuring **ethernet-switching** as the interface family on unit 0. The **ethernet-switching** family can be configured only on logical interface unit 0, and no other logical units can be defined on that interface.

Some switching features, however, cannot be configured on logical interface unit 0. Features such as Q-in-Q tunneling require the logical interface to transmit VLAN-tagged frames. To enable a logical interface to receive and forward Ethernet frames tagged with a matching VLAN ID, you must bind the logical interface to that VLAN. These features must be configured on a logical interface unit other than 0, because you cannot bind a VLAN ID to unit 0.

When you encapsulate an interface by using flexible Ethernet services, you can configure a logical interface unit other than 0 with **family ethernet-switching**. You can also configure other logical interfaces on that same interface with different types of Ethernet encapsulations. This enables logical interfaces that are bound to a VLAN ID to coexist with logical interfaces configured with **family ethernet-switching**.

For example, if you configure PVLAN on the same physical interface on which you are configuring Q-in-Q tunneling, you can use flexible ethernet services to support the enterprise style of configuration for PVLAN, using **family ethernet-switching**, along with **vlan-bridge** encapsulation for Q-in-Q tunneling.

BEST PRACTICE: We recommend you configure the following statements using groups when configuring devices that function as hardware VTEPs:

- set interfaces *interface-name* flexible-vlan-tagging
- set interfaces *interface-name* encapsulation extended-vlan-bridge
- set interfaces *interface-name* native-vlan-id *vlan-id*

To configure the interface to support both the service provider and enterprise styles of configuration:

1. Enable flexible Ethernet services encapsulation on the interface. The **flexible-ethernet-services** statement allows configuration of both service-provider-style logical interfaces and enterprise-style logical interfaces:

```
[edit interfaces interface-name]
```

```
user@switch# set encapsulation flexible-ethernet-services
```

2. Enable the interface to transmit packets with 802.1Q VLAN single-tagged and dual-tagged frames:

```
[edit interfaces interface-name]
```

```
user@switch# set flexible-vlan-tagging
```

3. Configure a logical interface (unit) on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# set unit unit-number
```

NOTE: Do not use logical interface unit 0. You must later bind a VLAN tag ID to the unit you specify in this step, and you cannot bind a VLAN tag ID to unit 0. It is a best practice to match the unit number to the VLAN ID to which the interface is bound.

4. Encapsulate the logical interface for service provider style bridging configuration—for example, use **vlan-bridge** encapsulation on an interface to be used for Q-in-Q tunneling. (If you were configuring the interface only for Q-in-Q tunneling, you would use **encapsulation extended-vlan-bridge** on the *physical* interface.)

```
[edit interfaces interface-name]
```

```
user@switch# set unit unit-number encapsulation vlan-bridge
```

5. Bind the logical interface from the preceding step to a VLAN ID:

```
[edit interfaces interface-name]
```

```
user@switch# set unit unit-number vlan-id vlan-id
```

6. Configure another logical interface. (If you were configuring just PVLAN, we would recommend that you configure a single logical interface for all PVLAN domains on an interface.)

```
[edit interfaces interface-name]
```

```
user@switch# set unit unit-number
```

7. Enable the logical interface in the preceding step for enterprise style bridging configuration:

[edit interfaces *interface-name*]

```
user@switch# set unit unit-number family ethernet-switching
```

8. Assign VLAN membership to the logical interface:

[edit interfaces *interface-name*]

```
user@switch# set unit unit-number family ethernet-switching vlan members vlan-id
```

9. Configure the interface as a trunk interface to transmit frames with 802.1Q VLAN tags:

[edit interfaces *interface-name*]

```
user@switch# set unit unit-number family ethernet-switching interface-mode trunk
```

NOTE: For EX4300 device, the service provider style configuration (encapsulation extended-vlan-bridge) is recommended only for QinQ scenarios. For other scenarios, use the enterprise style configuration.

SEE ALSO

Configuring Q-in-Q Tunneling on QFX Series, NFX Series, and EX4600 Switches with ELS Support
Creating a Private VLAN on a Single Switch with ELS Support (CLI Procedure)

4

CHAPTER

Monitoring and Troubleshooting Information

Monitoring Interfaces | **392**

Troubleshooting Interfaces | **399**

Monitoring Interfaces

IN THIS SECTION

- [Monitoring Interface Status and Traffic | 392](#)
- [Monitoring System Process Information | 392](#)
- [Monitoring System Properties | 393](#)
- [Monitor Statistics for a Fast Ethernet or Gigabit Ethernet Interface | 395](#)
- [Tracing Operations of the Interface Process | 397](#)

The below topics discuss the monitoring of the status and traffic, system process information, system properties, statistics for a fast Ethernet and the tracing operations of the interface process.

Monitoring Interface Status and Traffic

Purpose

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

Action

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

Meaning

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

Monitoring System Process Information

Purpose

View the processes running on the device.

Action

To view the software processes running on the device:

```
user@switch> show system processes
```

Meaning

[Table 61 on page 393](#) summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

Table 61: Summary of System Process Information Output Fields

Field	Values
PID	Identifier of the process.
Name	Owner of the process.
State	Current state of the process.
CPU Load	Percentage of the CPU that is being used by the process.
Memory Utilization	Amount of memory that is being used by the process.
Start Time	Time of day when the process started.

SEE ALSO

|

```
show system uptime
```

Monitoring System Properties

Purpose

View system properties such as the name, IP address, and resource usage.

Action

To monitor system properties in the CLI, enter the following commands:

- *show system uptime*
- *show system users*
- *show system storage*

Meaning

Table 62 on page 394 summarizes key output fields in the system properties display.

Table 62: Summary of Key System Properties Output Fields

Field	Values	Additional Information
General Information		
Serial Number	Serial number of device.	
Junos OS Version	Version of Junos OS active on the switch, including whether the software is for domestic or export use.	Export software is for use outside the USA and Canada.
Hostname	Name of the device.	
IP Address	IP address of the device.	
Loopback Address	Loopback address.	
Domain Name Server	Address of the domain name server.	
Time Zone	Time zone on the device.	
Time		
Current Time	Current system time, in Coordinated Universal Time (UTC).	
System Booted Time	Date and time when the device was last booted and how long it has been running.	
Protocol Started Time	Date and time when the protocols were last started and how long they have been running.	
Last Configured Time	Date and time when a configuration was last committed. This field also shows the name of the user who issued the last commit command.	
Load Average	CPU load average for 1, 5, and 15 minutes.	
Storage Media		

Table 62: Summary of Key System Properties Output Fields (continued)

Field	Values	Additional Information
Internal Flash Memory	Usage details of internal flash memory.	
External Flash Memory	Usage details of external USB flash memory.	
Logged in Users Details		
User	Username of any user logged in to the switch.	
Terminal	Terminal through which the user is logged in.	
From	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.	
Login Time	Time when the user logged in.	This is the user@switch field in show system users command output.
Idle Time	How long the user has been idle.	

SEE ALSO

| *show system processes*

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* for the fiber-optic Ethernet interface specifications.
3. Measure the transmit light level on the Tx port to verify that it is within specification. See *Fiber-Optic Ethernet Interface Specifications* for the optical specifications.

Tracing Operations of the Interface Process

To trace the operations of the router or switch interface process, dcd, perform the following steps:

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

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

2. Configure the **traceoptions** statement.

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

3. Configure the **no-remote-trace** option to disable remote tracing.

```
[edit interfaces traceoptions]
user@host# set no-remote-trace
```

4. Configure the **file filename** option.

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

5. Configure the **files number** option, **match regular-expression** option, **size size** option, and **world-readable | no-world-readable** option.

```
[edit interfaces traceoptions file]
user@host# set files number
user@host# set match regular-expression
user@host# set size size
user@host# set word-readable | no-world-readable
```

6. Configure the tracing flag.

```
[edit interfaces traceoptions]
user@host# set flag flag-option
```

7. Configure the **disable** option in **flag flag-option** statement to 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 **all**.

```
[edit interfaces traceoptions]
user@host# set flag flag-option disable
```

You can specify the following flags in the **interfaces traceoptions** statement:

- **all**—Enable all configuration logging.
- **change-events**—Log changes that produce configuration events.
- **gres-events**—Log the events related to GRES.
- **resource-usage**—Log the resource usage for different states.
- **config-states**—Log the configuration state machine changes.
- **kernel**—Log configuration IPC messages to kernel.
- **kernel-detail**—Log details of configuration messages to kernel.
- **select-events**—Log the events on select state machine.

By default, interface process operations are placed in the file named dcd and three 1-MB files of tracing information are maintained.

For general information about tracing, see the tracing and logging information in the *Junos OS Administration Library*.

SEE ALSO

Tracing Interface Operations Overview

Tracing Operations of an Individual Router Interface

[traceoptions](#) | [495](#)

Troubleshooting Interfaces

IN THIS SECTION

- [Troubleshooting Network Interfaces](#) | [399](#)
- [Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\)](#) | [400](#)
- [Troubleshooting Uplink Ports on EX2300 Switches](#) | [403](#)

The below topics discuss the troubleshooting of network interfaces and diagnosing a faulty twisted-pair cable.

Troubleshooting Network Interfaces

Statistics for logical interfaces on Layer 2 interfaces are not accurate

Problem

Description: On QFX5000 switches, statistics for logical interfaces are not supported on Layer 2 interfaces or on any child member interfaces of Layer 2 aggregated Ethernet (AE) interfaces—that is, output for the **show interfaces *interface-name*** operational-mode command does not provide accurate I/O information for the logical interfaces.

Solution

If you need to see statistics for those logical interfaces, configure firewall filter rules to collect the information.

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

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

Symptoms: When you check the status with the CLI command **show interfaces *interface-name***, the disabled port is not listed.

Cause

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

Solution

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

Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure)

Problem

Description: A 10/100/1000BASE-T Ethernet interface has connectivity problems that you suspect might be caused by a faulty cable.

Solution

Use the time domain reflectometry (TDR) test to determine whether a twisted-pair Ethernet cable is faulty.

The TDR test:

- Detects and reports faults for each twisted pair in an Ethernet cable. Faults detected include open circuits, short circuits, and impedance mismatches.
- Reports the distance to fault to within 1 meter.
- Detects and reports pair swaps, pair polarity reversals, and excessive pair skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX2300, EX3200, EX3300, EX3400, EX4200, and EX4300 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 network interfaces on line cards.

NOTE: We recommend running the TDR test on an interface when there is no traffic on the interface.

To diagnose a cable problem by running the TDR test:

1. Run the [request diagnostics tdr](#) command.

```
user@switch> request diagnostics tdr start interface ge-0/0/10

Interface TDR detail:
Test status           : Test successfully executed ge-0/0/10
```

2. View the results of the TDR test with the [show diagnostics tdr](#) command.

```
user@switch> show diagnostics tdr interface ge-0/0/10

Interface TDR detail:
Interface name           : ge-0/0/10
Test status              : Passed
Link status              : Down
MDI pair                 : 1-2
  Cable status            : Normal
  Distance fault          : 0 Meters
  Polartiy swap           : N/A
  Skew time               : N/A
MDI pair                 : 3-6
  Cable status            : Normal
  Distance fault          : 0 Meters
  Polartiy swap           : N/A
  Skew time               : N/A
MDI pair                 : 4-5
  Cable status            : Open
  Distance fault          : 1 Meters
  Polartiy swap           : N/A
  Skew time               : N/A
```

```

MDI pair           : 7-8
  Cable status      : Normal
  Distance fault    : 0 Meters
  Polartiy swap     : N/A
  Skew time         : N/A
Channel pair       : 1
  Pair swap         : N/A
Channel pair       : 2
  Pair swap         : N/A
Downshift          : N/A

```

3. Examine the **Cable status** field for the four MDI pairs to determine if the cable has a fault. In the preceding example, the twisted pair on pins 4 and 5 is broken or cut at approximately one meter from the **ge-0/0/10** port connection.

NOTE: The **Test Status** field indicates the status of the TDR test, not the cable. The value **Passed** means the test completed—it does not mean that the cable has no faults.

The following is additional information about the TDR test:

- The TDR test can take some seconds to complete. If the test is still running when you execute the **show diagnostics tdr** command, the **Test status** field displays **Started**. For example:

```

user@switch> show diagnostics tdr interface ge-0/0/22

Interface TDR detail:
Interface name           : ge-0/0/22
Test status              : Started

```

- You can terminate a running TDR test before it completes by using the **request diagnostics tdr abort interface interface-name** command. The test terminates with no results, and the results from any previous test are cleared.
- You can display summary information about the last TDR test results for all interfaces on the switch that support the TDR test by not specifying an interface name with the **show diagnostics tdr** command. For example:

```

user@switch> show diagnostics tdr

Interface   Test status   Link status   Cable status   Max distance fault

```

ge-0/0/0	Passed	UP	OK	0
ge-0/0/1	Not Started	N/A	N/A	N/A
ge-0/0/2	Passed	UP	OK	0
ge-0/0/3	Not Started	N/A	N/A	N/A
ge-0/0/4	Passed	UP	OK	0
ge-0/0/5	Passed	UP	OK	0
ge-0/0/6	Passed	UP	OK	0
ge-0/0/7	Not Started	N/A	N/A	N/A
ge-0/0/8	Passed	Down	OK	0
ge-0/0/9	Not Started	N/A	N/A	N/A
ge-0/0/10	Passed	Down	Fault	1
ge-0/0/11	Passed	UP	OK	0
ge-0/0/12	Not Started	N/A	N/A	N/A
ge-0/0/13	Not Started	N/A	N/A	N/A
ge-0/0/14	Not Started	N/A	N/A	N/A
ge-0/0/15	Not Started	N/A	N/A	N/A
ge-0/0/16	Not Started	N/A	N/A	N/A
ge-0/0/17	Not Started	N/A	N/A	N/A
ge-0/0/18	Not Started	N/A	N/A	N/A
ge-0/0/19	Passed	Down	OK	0
ge-0/0/20	Not Started	N/A	N/A	N/A
ge-0/0/21	Not Started	N/A	N/A	N/A
ge-0/0/22	Passed	UP	OK	0
ge-0/0/23	Not Started	N/A	N/A	N/A

SEE ALSO

Troubleshooting Interface Configuration and Cable Faults

[request diagnostics tdr | 806](#)

[show diagnostics tdr | 812](#)

Troubleshooting Uplink Ports on EX2300 Switches

This topic provides troubleshooting information for specific problems related to interfaces on EX2300 switches.


- [Speeds 10-Mbps and 100-Mbps not supported on uplink ports 4 and 5 on EX2300-48MP switches | 404](#)

Speeds 10-Mbps and 100-Mbps not supported on uplink ports 4 and 5 on EX2300-48MP switches

Problem

Description: The uplink ports 4 and 5 (see [Figure 13 on page 404](#)) do not support the speeds 10-Mbps and 100-Mbps.

Figure 13: Front Panel of an EX2300-48MP Switch



1—QR code	5—10-Gigabit Ethernet uplink ports
2—System LEDs	6—Management port
3—10/100/1000 BASE-T Gigabit Ethernet ports with PoE/PoE+ capability	7—Console port
4—100/1000/2500 BASE-T Gigabit Ethernet ports	8—USB port

Environment: A transceiver is installed in the uplink port 4 or 5 or both.

Symptoms: When you check the status with the CLI command [show interfaces ge](#) or with the J-Web user interface, the port is not listed.

Cause

EX2300-48MP switches do not support 10-Mbps and 100-Mbps speeds on uplink ports 4 and 5. This is an ASIC limitation.

Solution

Use the other ports if you need 10-Mbps and 100-Mbps speeds.

SEE ALSO

1

PART

Configuration Statements and Operational Commands

Configuration Statements: Interfaces | **407**

Configuration Statements: Gigabit Ethernet Interfaces | **518**

Configuration Statements: OTN Interfaces | **524**

Configuration Statements: Aggregated Ethernet Interfaces | **547**

Configuration Statements: Channelizing Interfaces | **648**

Configuration Statements: Energy Efficient Interfaces | **665**

Configuration Statements: VLANs | **666**

Configuration Statements: Uplink Failure Detection for Interfaces | **684**

Configuration Statements: Unicast Reverse Path Forwarding (uRPF) | **690**

Configuration Statements: IP Directed and Targeted Broadcast | **705**

Configuration Statements: ARP | **714**

Configuration Statements: Resilient Hashing | **723**

Configuration Statements: Generic Routing Encapsulation (GRE) | **744**

Configuration Statements: Flexible Ethernet Services Encapsulation | **764**

Operational Commands | **778**

Configuration Statements: Interfaces

IN THIS CHAPTER

- address | 409
- auto-negotiation | 413
- autostate-exclude | 414
- bandwidth (Interfaces) | 416
- broadcast | 418
- ccc | 419
- configured-flow-control | 420
- description (Interfaces) | 422
- disable (Interface) | 424
- ethernet (Alarm) | 426
- ethernet-switching | 427
- eui-64 | 428
- family | 429
- forward-and-send-to-re | 434
- forward-only | 435
- filter | 436
- hold-time (Physical Interface) | 438
- inet (interfaces) | 441
- inet6 (interfaces) | 442
- inet (enhanced-hash-key) | 443
- inet6 (enhanced-hash-key) | 446
- interface (Multichassis Protection) | 448
- interface-mode | 449
- interface-range | 451
- interfaces (QFX Series) | 453
- interfaces (EX Series switches) | 464
- irb (Interfaces) | 475
- loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | 480

- [mac](#) | 482
- [media-type \(Dual-Purpose Uplink Ports\)](#) | 483
- [member](#) | 484
- [member-range](#) | 485
- [mode \(Interfaces\)](#) | 486
- [mtu](#) | 487
- [nd6-stale-time](#) | 489
- [no-redirects](#) | 490
- [policer \(MAC\)](#) | 491
- [preferred](#) | 493
- [primary \(Address on Interface\)](#) | 494
- [traceoptions \(Individual Interfaces\)](#) | 495
- [reflective-relay](#) | 504
- [speed \(Ethernet\)](#) | 505
- [traps](#) | 513
- [unidirectional](#) | 515
- [unit](#) | 516

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” on page 76](#).

- 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>
<i>negotiate-address</i>
<i>unnumbered-address (Ethernet)</i>

auto-negotiation

Syntax

```
(auto-negotiation | no-auto-negotiation);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Explicitly enable or disable autonegotiation. Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 1G and 10G.

- **auto-negotiation**—Enable autonegotiation.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure link mode and speed options.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[speed](#) | [505](#)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | [109](#)

Junos OS Network Interfaces Library for Routing Devices

autostate-exclude

Syntax

```
autostate-exclude;
```

Hierarchy Level

```
[edit interface interface-name ether-options]
```

Release Information

Statement introduced in Junos OS Release 14.1x53-D40 and Junos OS Release 17.3R1 on QFX5100 switches.

Description

Specify not to include an IRB interface in the state calculation for VLAN members. The default behavior is not to exclude an IRB interface in the state calculation unless all the ports on the interface go down. Because an IRB interface often has multiple ports in a single VLAN, the state calculation for a VLAN member might include a port that is down, possibly resulting in traffic loss. This feature enables you to exclude a trunk or access interface from the state calculation, which results in the IRB interface being marked as down as soon as the port specifically assigned to a VLAN goes down.

IRB interfaces are used to bind specific VLANs to Layer 3 interfaces, enabling a switch to forward packets between those VLANs— without having to configure another device, such as a router, to connect VLANs. In a typical scenario, a port on the interface is assigned to a specific VLAN, while a different port on that interface is assigned to an 802.1Q trunk interface to carry traffic between multiple VLANs, and a third port on that interface is assigned to an access interface used to connect the VLAN to network devices.

To ensure that an interface is marked as down and thereby excluded from the state calculation for VLAN members when the port assigned to the VLAN goes down, configure this statement on the trunk or access interface. The trunk or port interface is automatically excluded from the state calculation of the IRB interface. In this way, when a port assigned to a specified VLAN goes down, the IRB interface assigned to that VLAN is also marked as down.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration

RELATED DOCUMENTATION

Excluding an IRB Interface from State Calculations on a QFX Series Switch

[port-mode](#) | 676

| *show ethernet-switching interface*

bandwidth (Interfaces)

Syntax

```
bandwidth rate;
```

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 9.0 for EX Series switches.

Description

Configure the bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.

NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

$$\text{cost} = \text{reference-bandwidth} / \text{bandwidth},$$

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

Options

rate—Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value 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). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation **c**; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.

Range: Not limited.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Interface Bandwidth](#) | 81

broadcast

Syntax

```
broadcast address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family address address],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family address  
  address]
```

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

Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0 (0.0.0.0), nor can you specify a broadcast address (255.255.255.255). For example, in the statement **set interface ge-0/0/0 unit 0 family inet address 10.1.1.0/24**, the subnet address **10.1.1.0** has the host address of 0. Hence, you cannot configure this address. Similarly, for the subnet, you cannot use the broadcast address **10.1.1.255/24**.

Default

The default broadcast address has a host portion of all ones.

Options

address—Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses **0.0.0.0** or **255.255.255.255**.

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

CCC

Syntax

```
ccc;
```

Hierarchy Level

```
[edit interfaces ge-fpc/slot/ port unit logical-unit-number family]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure the logical interface as a circuit cross-connect (CCC).

NOTE: On QFX10002 switches, circuit cross-connects are not supported on aggregated Ethernet interfaces.

Default

You must configure a logical interface to be able to use the physical device.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

configured-flow-control

Syntax

```
configured-flow-control {  
    rx-buffers (on | off);  
    tx-buffers (on | off);  
}
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Release Information

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the **rx-buffers** and the **tx-buffers** values when you configure asymmetric flow control.

Use the **flow-control** and **no-flow-control** statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.

NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.

Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

Default

Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.

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

<i>congestion-notification-profile</i>
flow-control 589
<i>Configuring CoS Asymmetric Ethernet PAUSE Flow Control</i>
<i>Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control</i>
<i>Understanding CoS Flow Control (Ethernet PAUSE and PFC)</i>

description (Interfaces)

Syntax

```
description text;
```

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.

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 14.1X53-D20 for the OCX Series.

Description

Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.

The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.

Options

text—Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Interface Description

[Adding a Logical Unit Description to the Configuration | 78](#)

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches | 119](#)

[Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support | 114](#)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

Using DHCP Relay Agent Option 82 Information

Junos OS Network Interfaces Library for Routing Devices

Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support

disable (Interface)

Syntax

```
disable;
```

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.

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

Disable a physical or a logical interface, effectively unconfiguring it.



CAUTION:

- Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.
- If aggregated SONET links are configured between a T1600 router and a T4000 router, interface traffic is disrupted when you disable the physical interface configured on the T1600 router. If you want to remove the interface, we recommend that you deactivate the interface instead of disabling it.

NOTE:

- When you use the **disable** statement at the **[edit interfaces]** hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet (GE) PICs with SFP and XFP transceivers and ATM MIC with SFP do support it and the laser will be turned off when the interface is disabled. If the ATM MIC with SFP is part of an APS group, then the laser will not be turned off when you use the **disable** statement at the **[edit interfaces]** hierarchy level..
- When you disable or deactivate an interface, then all the references made to the deactivated interface must be removed from the routing instance.
- For abstracted fabric interfaces, the **disable** command disables AF interface on the local GNF only.



WARNING: Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling a Physical Interface | 72](#)

[Disabling a Logical Interface | 85](#)

ethernet (Alarm)

Syntax

```
ethernet {  
    link-down (red | yellow | ignore);  
}
```

Hierarchy Level

```
[edit chassis alarm],  
[edit chassis interconnect-device name alarm],  
[edit chassis node-group name alarm]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure alarms for an Ethernet interface.

Options

The remaining 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.

ethernet-switching

Syntax

```
ethernet-switching {
  filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
  }
  interface-mode (access | trunk);
  recovery-timeout seconds;
  storm-control profile-name;
  vlan {
    members (vlan-name [[-vlan-names] | all]);
  }
}
```

Hierarchy Level

[edit **interfaces** *ge-chassis/slot/port unit logical-unit-number*] family

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure Ethernet switching protocol family information for the logical interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Default

You must configure a logical interface to be able to use the physical device.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

eui-64

Syntax

```
eui-64;
```

Hierarchy Level

```
[edit interfaces interface-name unit number family inet6 address address]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Interface Address](#) | 76

family

Syntax

The **ethernet-switching** statement and all of its substatements are not supported on OCX Series switches.

```
family {
  ethernet-switching {
    filter {
      group filter-group-number;
      input filter-name;
      input-list [ filter-names ];
      output filter-name;
      output-list [ filter-names ];
    }
    interface-mode (access | trunk);
    recovery-timeout seconds;
    storm-control profile-name;
    vlan {
      members (vlan-name |[-vlan-names] | all);
    }
  }
  fibre-channel {
    port-mode (f-port | np-port);
  }
  inet {
    accounting {
      destination-class-usage;
      source-class-usage {
        input;
        output;
      }
    }
  }
  address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-number {
      (accept-data | no-accept-data);
      advertise-interval seconds;
      advertisements-threshold number;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
    }
  }
}
```

```

    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [addresses];
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}

```

```

inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-name;
            active-interface interface-name;
        }
    }
}
(dad-disable | no-dad-disable);
filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
}

```

```

        output filter-name;
        output-list [ filter-names ];
    }
    mtu bytes;
    nd6-stale-time time;
    no-neighbor-learn;
    no-redirects;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
    mpls {
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        mtu bytes;
    }
}
}

```

Hierarchy Level

```

[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range interface-name unit logical-unit-number family]

```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure protocol family information for the logical interface on the QFX Series and OCX Series product.

Default

NOTE: The **ethernet-switching** statement and all of its substatements are not supported on OCX Series switches.

Access interfaces on the QFX Series are set to **family ethernet-switching** by default. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting first.

You must configure a logical interface to be able to use the physical device.

Options

Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)

NOTE: Routed VLAN interfaces, also referred to as integrated routing and bridging (IRB) interfaces, are not supported on OCX Series switches.

- 10-Gigabit Ethernet (**xe**)

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109](#)

[Configuring Link Aggregation | 281](#)

[Configuring IRB Interfaces on Switches](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

forward-and-send-to-re

Syntax

```
forward-and-send-to-re;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet targeted-broadcast],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet  
targeted-broadcast]
```

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface and the Routing Engine. The packets are broadcast only if the egress interface is a LAN interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Targeted Broadcast | 229](#)

[targeted-broadcast | 706](#)

[Understanding Targeted Broadcast | 226](#)

forward-only

Syntax

```
forward-only;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet targeted-broadcast],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet  
targeted-broadcast]
```

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Targeted Broadcast | 229](#)

[targeted-broadcast | 706](#)

[Understanding Targeted Broadcast | 226](#)

filter

Syntax

```
filter {
  group filter-group-number;
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
```

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 9.0 for EX Series switches.

Description

NOTE: On EX Series switches, the **group**, **input-list**, **output-filter** statements are not supported under the [edit interfaces *interface-name* unit *logical-unit-number* family *inet*], [edit interfaces *interface-name* unit *logical-unit-number* family *inet6*], and [edit interfaces *interface-name* unit *logical-unit-number* family *mpls*] hierarchies.

Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the **family ethernet-switching**, **inet**, **inet6**, **mpls**, or **vpls** only.

Options

group filter-group-number—Define an interface to be part of a filter group. The default filter group number is 0.

Range: 0 through 255

input filter-name—Name of one filter to evaluate when packets are received on the interface.

output filter-name—Name of one filter to evaluate when packets are transmitted on the 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

<i>Applying a Filter to an Interface</i>
<i>Junos OS Administration Library</i>
<i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>
<i>Configuring Firewall Filters (CLI Procedure)</i>
<i>family</i>

hold-time (Physical Interface)

Syntax

```
hold-time up milliseconds down milliseconds;
```

Hierarchy Level

```
[edit interfaces interface-name],  
[edit interfaces interface-range interface-range-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 10.4R5 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Statement introduced in Junos OS Release 12.1 for the SRX Series.

Description

Specify the **hold-time** value to use to damp shorter interface transitions milliseconds. The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still down, then the router begins to advertise the interface as being down. Similarly, when a hold-up timer is configured and an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still up, then the router begins to advertise the interface as being up.

We recommend you to configure sub-second down hold timer value as 500 milliseconds on:

- MX204 router
- MX routers with MPC7E, MPC8E, MPC9E, MPC10E, and MX10003-LC2103 MPCs

The sub-second down hold timer is not supported on the interfaces created on MX routers with MPC3E, MPC4E, MPC5E, MPC5EQ, MPC6E, and MPC11E MPCs. The sub-second hold timer is a configuration of hold timer less than a second.

NOTE:

- We recommend that you configure the **hold-time** value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for **hold-time** depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.
- The **hold-time** option is not available for controller interfaces.

NOTE: On MX Series routers with MPC3E and MPC4E, we recommend that you do not configure the hold-down timer to be less than 1 second. On MX Series routers with MPC5EQ-100G10G (MPC5EQ) or MPC6E (MX2K-MPC6E) with 100-Gigabit Ethernet MIC with CFP2 OTN interfaces, we recommend that you do not configure the hold-down timer to be less than 3 seconds.

Default

Interface transitions are not damped.

Options

down milliseconds—Hold time to use when an interface transitions from up to down. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

up milliseconds—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>advertise-interval</i>
interfaces (EX Series switches) 464
<i>Physical Interface Damping Overview</i>
Damping Shorter Physical Interface Transitions 59
<i>Damping Longer Physical Interface Transitions</i>

inet (interfaces)

Syntax

```
inet {  
  address address {  
    primary;  
    filter input filter-name;  
    filter output filter-name;  
    targeted-broadcast;  
  }  
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family],  
[edit interfaces interface-range interface-name unit logical-unit-number family]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the primary IP address for the logical interface.

Default

You must configure a logical interface to be able to use the physical device.

Options

The remaining statements are explained separately. 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 and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

inet6 (interfaces)

Syntax

```
inet6 {
  address address {
    eui-64
    preferred
    primary;
    filter input filter-name;
    filter output filter-name;
  }
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family],
[edit interfaces interface-range interface-name unit logical-unit-number family]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the primary IP address for the logical interface.

Default

You must configure a logical interface to be able to use the physical device.

Options

The remaining statements are explained separately. 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 and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches](#) | 119

inet (enhanced-hash-key)

List of Syntax

[Syntax \(EX Series and QFX5100 Switch\) on page 443](#)

[Syntax \(QFX10000 Series Switches\) on page 443](#)

Syntax (EX Series and QFX5100 Switch)

```
inet {
  gtp-tunnel-endpoint-identifier;
  no-ipv4-destination-address;
  no-ipv4-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-protocol;
  vlan-id;
}
```

Syntax (QFX10000 Series Switches)

```
inet {
  gtp-tunnel-endpoint-identifier;
  no-ipv4-destination-address;
  no-ipv4-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-incoming-port;
}
```

Hierarchy Level

[edit forwarding-options [enhanced-hash-key](#) family]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

Statement introduced in Junos OS Release 15.1X53-D30 on QFX10000 Series Switches.

Description

Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.

When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the **inet** statement and uses the information in the fields to decide how to place

traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to **layer2-payload**. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the **set forwarding-options enhanced-hash-key hash-mode layer2-payload** statement.

Default

The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:

- IP destination address
- IP source address
- Layer 4 destination port
- Layer 4 source port
- Protocol

Options

no-ipv4-destination-address—Exclude the IPv4 destination address field from the hashing algorithm.

no-ipv4-source-address—Exclude the IPv4 source address field from the hashing algorithm.

no-l4-destination-port—Exclude the Layer 4 destination port field from the hashing algorithm.

no-l4-source-port—Exclude the Layer 4 source port field from the hashing algorithm.

no-protocol—Exclude the protocol field from the hashing algorithm.

no-incoming-port—Exclude the incoming port number from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

NOTE: The **vlan-id** option is not supported and should not be configured on a Virtual Chassis or Virtual Chassis Fabric (VCF) that contains any of the following switches as members: EX4300, EX4600, QFX3500, QFX3600, QFX5100, or QFX5110 switches.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)	349
Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic	342
<i>Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)</i>	
Understanding Per-Packet Load Balancing	257
hash-seed	735
enhanced-hash-key	725
hash-mode	733
inet6	446

inet6 (enhanced-hash-key)

List of Syntax

[Syntax \(EX Series and QFX5100 Switch\) on page 446](#)

[Syntax \(QFX10000 Series Switches\) on page 446](#)

Syntax (EX Series and QFX5100 Switch)

```
inet6 {
  no-ipv6-destination-address;
  no-ipv6-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-next-header;
  vlan-id;
}
```

Syntax (QFX10000 Series Switches)

```
inet6 {
  gtp-tunnel-endpoint-identifier;
  ipv6-flow-label;
  no-ipv6-destination-address;
  no-ipv6-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-incoming-port;
}
```

Hierarchy Level

[edit forwarding-options [enhanced-hash-key](#) family]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 on EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 on QFX Series devices.

Statement introduced in Junos OS Release 15.1X53-D30 on QFX10000 Series switches.

Description

Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.

When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place

traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the **set forwarding-options enhanced-hash-key hash-mode layer2-payload** statement.

Default

The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:

- IP destination address
- IP source address
- Layer 4 destination port
- Layer 4 source port
- Next header

Options

no-ipv6-destination-address—Exclude the IPv6 destination address field from the hashing algorithm.

no-ipv6-source-address—Exclude the IPv6 source address field from the hashing algorithm.

no-l4-destination-port—Exclude the Layer 4 destination port field from the hashing algorithm.

no-l4-source-port—Exclude the Layer 4 source port field from the hashing algorithm.

no-incoming-port—Exclude the incoming port number from the hashing algorithm.

no-next-header—Exclude the Next Header field from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) | 349](#)

[Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic | 342](#)

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

[Understanding Per-Packet Load Balancing | 257](#)

[hash-seed | 735](#)

[enhanced-hash-key | 725](#)

[hash-mode | 733](#)

[inet | 443](#)

interface (Multichassis Protection)

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit multi-chassis multi-chassis-protection peer]
```

Release Information

Statement introduced in Junos OS Release 9.6 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the name of the interface that is being used as an interchassis link-protection link (ICL-PL). The two switches hosting a multichassis link aggregation group (MC-LAG) use this link to pass Inter-Chassis Control Protocol (ICCP) and data traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

interface-mode

Syntax

```
interface-mode (access | trunk <inter-switch-link>);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family bridge],  
[edit interfaces interface-name unit logical-unit-number family ethernet-switching],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family bridge]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 15.1.

inter-switch-link option introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description

NOTE: This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [port-mode](#). For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

QFX3500 and QFX3600 standalone switches—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.

NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see *Configuring a Trunk Interface on a Bridge Network*.

Options

access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the **vlan-id** statement.

trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the **vlan-id** or **vlan-id-list** statement.

trunk inter-switch-link—For a private VLAN, configure the InterSwitch Link protocol (ISL) on a trunk port of the primary VLAN in order to connect the switches composing the PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Access Mode on a Logical Interface

Configuring a Logical Interface for Trunk Mode

Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support

Tunnel Services Overview

Tunnel Interface Configuration on MX Series Routers Overview

interface-range

Syntax

The **vlan-id** statement is not supported on OCX Series switches.

```
interface-range interface-range-name {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
  }
  (auto-negotiation| no-auto-negotiation);
  (flow-control | no-flow-control);
  link-mode mode;
  speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
  description text;
  disable;
  family family-name {...}
  (traps | no traps);
  vlan-id vlan-id-number;
}
}
```

Hierarchy Level

[edit [interfaces](#)]

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

NOTE: The **vlan-id** statement and Fibre Channel interfaces are not supported on OCX Series switches.

NOTE: The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces. Interface ranges are not supported on channelized interfaces.

Group interfaces that share a common configuration profile.

Options

interface-range-name—Name of the interface range.

NOTE: You can use regular expressions and wildcards to specify the interfaces in the member range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately. 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

[Understanding Interface Ranges for Switches | 87](#)

[Interfaces Overview for Switches | 30](#)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109](#)

Junos OS Network Interfaces Library for Routing Devices

interfaces (QFX Series)

Syntax

The following statements and their associated substatements are not supported on OCX Series switches: **auto-negotiation**, **speed**, **ethernet-switching**, **fcoe-lag**, **fibre-channel**, **fibrechannel-options**, **mc-ae**, **vlan**, **vlan-id**, and **vlan-tagging**.

```

interfaces {
  aex {
    disable;
    aggregated-ether-options {
      configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
      }
      (fcoe-lag | no-fcoe-lag);
      flexible-vlan-tagging;
      (flow-control | no-flow-control);
      lacp mode {
        admin-key key;
        force-up;
        periodic interval;
        system-id mac-address;
      }
      link-speed speed;
      local-bias;
      loopback;
      no-loopback;
      minimum-links number;
    }
    mc-ae {
      chassis-id chassis-id;
      mc-ae-id mc-ae-id;
      mode (active-active);
      status-control (active | standby);
    }
    description text;
    gratuitous-arp-reply | no-gratuitous-arp-reply
    hold-time down milliseconds up milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no traps);
    unit logical-unit-number {

```

```

disable;
description text;
family {
    ethernet-switching {
        filter input filter-name;
        filter output filter-name;
        native-vlan-id vlan-id;
        port-mode mode;
        reflective-relay;
        vlan {
            members [ (all | names | vlan-ids) ];
        }
    }
    inet {
        address address {
            primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
    }
    (traps | no traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}

```

```

interface-range interface-range-name {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
  }
  (auto-negotiation | no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  (flow-control | no-flow-control);
  link-mode mode;
  speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
  disable;
  description text;
  family family-name {...}
  (traps | no traps);
  vlan-id vlan-id-number;
}
}

```

```
lo0 {  
  disable;  
  description text;  
  hold-time milliseconds down milliseconds;  
  traceoptions;  
  (traps | no traps);  
  unit logical-unit-number {  
    disable;  
    description text;  
    family {  
      inet {  
        address address {  
          primary;  
        }  
        filter input filter-name;  
        filter output filter-name;  
        primary;  
        targeted-broadcast;  
      }  
    }  
    (traps | no traps);  
  }  
}
```



```

mex {
  disable;
  description text;
  hold-time milliseconds down milliseconds;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  no-gratuitous-arp-request;
  traceoptions;
  traps;
  unit logical-unit-number {
    disable;
    description text;
    family {
      ethernet-switching {
        filter input filter-name;
        filter output filter-name;
        native-vlan-id vlan-id;
        port-mode mode;
        reflective-relay;
        vlan {
          members [ (all | names | vlan-ids) ];
        }
      }
      inet {
        address address {
          primary;
          filter input filter-name;
          filter output filter-name;
          primary;
          targeted-broadcast;
        }
      }
    }
    traps;
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}

```

```

vlan {
  disable;
  description text;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no traps);
  unit logical-unit-number {
    description text;
    disable;
    family {
      inet {
        address address {
          primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
      }
      (traps | no traps);
    }
  }
}
fc-0/0/port {
  fibrechannel-options {
    bb-sc-n;
    (loopback | no-loopback);
    speed (auto-negotiation | 2g | 4g | 8g);
  }
  unit logical-unit-number {
    disable;
    description text;
    family {
      fibre-channel {
        port-mode np-port;
      }
      (traps | no traps);
    }
  }
}

```

```

ge-0/0/port {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
        primary;
      }
    }
  }
  (auto-negotiation | no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  flexible-vlan-tagging;
  (flow-control | no-flow-control);
  link-mode mode;
  loopback;
  no-loopback;
  speed (auto-negotiation | speed);
}
gratuitous-arp-reply| no-gratuitous-arp-reply);
hold-time milliseconds down milliseconds;
mac
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
  description text;
  disable;
  family {
    ethernet-switching {
      filter input filter-name;
      filter output filter-name;
      native-vlan-id vlan-id;
      port-mode mode;
      reflective-relay;
      vlan {
        members [ (all | names | vlan-ids) ];
      }
    }
  }
  inet {

```

```

        address address {
            primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
    }
    (traps | no traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
vrrp-group group-id {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
        }
        priority-hold-time seconds;
        route prefix/prefix-length routing-instance instance-name priority-cost priority;
    }
}
virtual-address [ addresses ];
}

```

```

xe-0/0/port {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
        (primary | backup);
      }
    }
    configured-flow-control {
      rx-buffers (on | off);
      tx-buffers (on | off);
    }
    flexible-vlan-tagging;
    (flow-control | no-flow-control);
    loopback;
    no-loopback;
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply)
  hold-time milliseconds down milliseconds;
  mac
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no traps);
  unit logical-unit-number {
    disable;
    description text;
    family {
      ethernet-switching {
        filter input filter-name;
        filter output filter-name;
        native-vlan-id vlan-id;
        port-mode mode;
        reflective-relay;
        vlan {
          members [ (all | names | vlan-ids) ];
        }
      }
      fibre-channel {
        port-mode (f-port | np-port);
      }
      inet {

```

```

    address address {
        primary;
    }
    filter input filter-name;
    filter output filter-name;
    primary;
    targeted-broadcast;
}
(traps | no traps);
vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the interfaces on the QFX Series and OCX Series.

The following statements and their associated substatements are not supported on OCX Series switches: **auto-negotiation**, **ethernet-switching**, **fcoe-lag**, **fibre-channel**, **fibrechannel-options**, **mc-ae**, **speed**, **vlan**, **vlan-id**, and **vlan-tagging**

Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.

Options

aex—Configure an aggregated Ethernet interface.

xe-0/0/port/—Configure a 10-Gigabit Ethernet interface.

ge-0/0/port/—Configure a Gigabit Ethernet interface.

fc-0/0/port/—Configure a Fibre Channel interface.

meX/—Configure a management interface.

mc-ae—Configure a multichassis aggregated Ethernet (MC-AE) 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

[Interfaces Overview for Switches | 30](#)

[Understanding Interface Ranges for Switches | 87](#)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109](#)

[Configuring Link Aggregation | 281](#)

Configuring a Layer 3 Logical Interface

interfaces (EX Series switches)

Syntax

[interfaces ae on page 464](#)
[interfaces ge on page 466](#)
[interfaces interface-range on page 467](#)
[interfaces lo0 on page 468](#)
[interfaces me0 on page 469](#)
[interfaces traceoptions on page 469](#)
[interfaces vlan on page 470](#)
[interfaces vme on page 471](#)
[interfaces xe on page 472](#)

interfaces ae

```

aex {
  accounting-profile name;
  aggregated-ether-options {
    (flow-control | no-flow-control);
    lACP {
      (active | passive);
      admin-key key;
      periodic interval;
      system-id mac-address;
    }
    (link-protection | no-link-protection);
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
  }
  description text;
  disable;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
  }
}

```



```
    bandwidth rate;  
    description text;  
    disable;  
    family family-name {...}  
    proxy-arp (restricted | unrestricted);  
    (traps | no-traps);  
    vlan-id vlan-id-number;  
  }  
  vlan-tagging;  
}
```

interfaces ge

```

ge-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lacp {
                force-up;
            }
        }
        (auto-negotiation | no-auto-negotiation);
        (flow-control | no-flow-control);
        ieee-802-3az-eee;
        link-mode mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    media-type;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

```

interfaces interface-range

```

interface-range name {
    accounting-profile name;
    description text;
    disable;
    ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lacp {
                force-up;
            }
        }
        (auto-negotiation | no-auto-negotiation);
        (flow-control | no-flow-control);
        ieee-802-3az-eee;
        link-mode mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    member interface-name;
    member-range starting-interface name to ending-interface name;
    mtu bytes;
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

```

interfaces lo0

```

lo0 {
    accounting-profile name;
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        (traps | no-traps);
    }
}

```

interfaces me0

```

me0 {
  accounting-profile name;
  description text;
  disable;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  no-gratuitous-arp-request;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}

```

interfaces traceoptions

```

traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}

```

interfaces vlan

```

vlan {
  accounting-profile name;
  description text;
  disable;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
  }
}

```

interfaces vme

```

vme {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

```

interfaces xe

```

xe-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    ether-options {
        802.3ad {
            aex;
            (backup | primary);
            lACP {
                force-up;
            }
        }
        (flow-control | no-flow-control);
        link-mode mode;
        (loopback | no-loopback);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}

```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure interfaces on EX Series switches.

Options

See [Table 63 on page 473](#) for the interface types and protocol-family options supported on the switch. Different protocol families support different subsets of the interface types on the switch. See the **family** statement for syntax of the protocol families supported for switch interfaces.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Table 63: Interface Types and Their Supported Protocol Families

Interface Type	Description	Supported Protocol Families					
		ccc	ethernet-switching	inet	inet6	iso	mpls
ae	Aggregated Ethernet interface (also referred to as a link aggregation group [LAG])	✓*	✓	✓	✓	✓	✓
ge	Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
interface-range	Interface-range configuration	Supported protocol families are the ones supported by the interface types that compose the range.					
lo0	Loopback interface			✓	✓	✓	✓
me0	Management Ethernet interface		✓	✓	✓	✓	✓
vlan	Routed VLAN interface (RVI)			✓	✓	✓	
vme	Virtual management Ethernet interface			✓	✓	✓	✓
xe	10-Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
		*Supported on EX8200 switches only					

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

<i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>
<i>Configuring Aggregated Ethernet Links (CLI Procedure)</i>
<i>Configuring a Layer 3 Subinterface (CLI Procedure)</i>
<i>Configuring Routed VLAN Interfaces on Switches (CLI Procedure)</i>
<i>Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)</i>
Interfaces Overview for Switches 30
Junos OS Interfaces Fundamentals Configuration Guide
Junos OS Ethernet Interfaces Configuration Guide

irb (Interfaces)

Syntax

```

irb {
    accounting-profile name;
    arp-l2-validate;
    description text;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    mtu bytes;
    no-gratuitous-arp-request;

    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        enhanced-convergence;
        disable;
        encapsulation type;
        family inet {
            accounting {
                destination-class-usage;
                source-class-usage {
                    input;
                    output;
                }
            }
        }
        address ipv4-address {
            arp ip-address (mac | multicast-mac) mac-address <publish>;
            broadcast address;
            preferred;
            primary;
            vrrp-group group-number {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                advertisements-threshold number;
                authentication-key key;
                authentication-type authentication;
                fast-interval milliseconds;
            }
        }
    }
}

```

```

(preempt | no-preempt) {
    hold-time seconds;
}
priority number;
track {
    interface interface-name {
        bandwidth-threshold bandwidth;
        priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-address [ addresses ];
vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
}
}
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}

```

```

family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-inet6-address [addresses];
        virtual-link-local-address ipv6-address;
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}

```

```

    }
    mtu bytes;
    nd6-stale-time seconds;
    no-neighbor-learn;
    no-redirects;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
no-auto-virtual-gateway-esi;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}

```

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

irb option introduced in Junos OS Release 13.2 for the QFX Series.

Description

Configure the properties of a specific integrated bridging and routing (IRB) interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

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

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

Configuring a Pseudowire Subscriber Logical Interface Device

media-type (Dual-Purpose Uplink Ports)

Syntax

```
media-type (copper | fiber);
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 11.3 for EX Series switches.

Statement introduced in Junos OS Release 12.2 for ACX1000 Universal Metro Routers.

Description

(EX2200-C switch only) Configure the media type for a dual-purpose uplink port (one RJ-45 port and one SFP port) on an EX2200 switch. If you use the media-type for a dual-purpose uplink port, the alternate media type cannot be used with the port.

(ACX1000 routers only) Select the media type (copper or fiber) for the 1-Gigabit Ethernet interfaces.

Default

When **media-type** is not set, the port accepts either type of connection. The default media type is fiber.

Options

copper—The dual-purpose uplink port accepts only a 10/100/1000BASE-T copper connection.

fiber—The dual-purpose uplink port accepts only an SFP fiber connection.

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 Type on Dual-Purpose Uplink Ports](#) | 71

member

Syntax

```
member interface-name;
```

Hierarchy Level

```
[edit interfaces interface-range interface-range-name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the name of the member interface belonging to an interface range on the QFX Series switch.

Options

interface-name—Name of the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109](#)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches | 119](#)

[Interfaces Overview for Switches | 30](#)

Junos OS Network Interfaces Library for Routing Devices

member-range

Syntax

```
member-range starting-interface-name ending-interface-name;
```

Hierarchy Level

```
[edit interfaces interface-range interface-range-name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.

Options

starting interface-name ending interface-name—Name of the first member and the name of the last member in the interface sequence.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Interface Ranges for Switches | 87](#)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109](#)

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches | 119](#)

[Interfaces Overview for Switches | 30](#)

[Interfaces Overview for Switches | 30](#)

Junos OS Network Interfaces Library for Routing Devices

mode (Interfaces)

Syntax

```
mode (strict | loose);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family (inet | inet6) rpf-check],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family (inet | inet6)  
rpf-check]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 15.1F6 for PTX Series routers with third-generations FPCs installed.

Description

Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.

In strict mode, unicast RPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix. If the incoming packet fails the unicast RPF check, the packet is not accepted on the interface.

Default

If you do not include this statement, unicast RPF is in strict mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Unicast RPF Strict Mode](#)

mtu

Syntax

```
mtu bytes;
```

Hierarchy Level

```
[edit interfaces interface-name],  
[edit interfaces interface-range interface-name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On QFX3500, QFX3600, QFX5100, and OCX Series switches, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight.

On QFX5110 switches, setting MTU on the L3 interface does not take effect and packets with MTU greater than the default value are dropped.

Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:

- **For LAG interfaces**—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (**aex**) automatically configures the jumbo MTU size on the member links.
- **For RVIs**—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named **vlan**. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the **vlan** interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named **vlan** (the RVI). On a QFX5100 switch jumbo frames on the RVI are configured on the basis of the interface MTU.

NOTE: RVIs are not supported on OCX Series switches.



CAUTION: Setting or deleting the jumbo MTU size on the RVI (the **vlan** interface) while the switch is transmitting packets might result in dropped packets.

Options

bytes —MTU size.

Range: 64 through 9216 bytes

Default: 1514 bytes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

Junos OS Network Interfaces Library for Routing Devices

nd6-stale-time

Syntax

```
nd6-stale-time seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Set the stale timer for IPv6 neighbor reachability confirmation. Reachability of the IPv6 neighbors is confirmed only after the stale timer has expired. For example, by setting the stale timer to 180 seconds, users can specify that IPv6 neighbor reachability be confirmed every 180 seconds.

NOTE: When the Routing Engine sends a control packet to an IPv6 neighbor, the stale timer is the maximum interval in which neighbor reachability is confirmed. In such cases, IPv6 neighbor reachability is confirmed before the stale timer expires.

Default

Default is 20 minutes (1200 seconds)

Options

seconds—Duration in seconds.

Range: 1 to 18000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

IPv6 Neighbor Discovery Overview

show ipv6 neighbors

no-redirects

Syntax

```
no-redirects;
```

Hierarchy Level

```
[edit 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 9.0 for EX Series switches.

Description

Do not send protocol redirect messages on the interface.

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the **[edit system]** hierarchy level.

Default

Interfaces send protocol redirect messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Disabling the Transmission of Redirect Messages on an Interface

Junos OS Administration Library

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*

preferred

Syntax

```
preferred;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family address address],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family address  
  address]
```

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

Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.

NOTE: The **edit logical-systems** hierarchy is not available on QFabric systems.

Default

The lowest-numbered address on the subnet is the preferred address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Interface Address](#) | 76

primary (Address on Interface)

Syntax

```
primary;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family address address],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family address  
  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 this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.

NOTE: The **edit logical-systems** hierarchy is not available on QFabric systems.

Default

For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Interface Address](#) | 76

traceoptions (Individual Interfaces)

List of Syntax

[Syntax \(Individual interfaces with PTX Series, EX Series, ACX Series\) on page 495](#)

[Syntax \(Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series\) on page 495](#)

[Syntax \(OAMLFM with EX Series, QFX Series, NFX Series\) on page 495](#)

[Syntax \(Interface process with ACX Series, SRX Series, MX Series, M Series, T Series\) on page 495](#)

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 64 on page 498](#) lists options for traceoption command for the following platforms:

Table 64: 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
file filename	<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 dcd.</p>
files number	<p>—(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.</p>		<p>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum xk to specify KB, xm to specify MB, or xg to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p>	<p>—(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.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p><i>Range: 2 through</i></p>

Table 64: 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 <i>Default: 3 files</i>
flag	<p>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"> • all—All interface tracing operations • event—Interface events • ipc—Interface interprocess communication (IPC) messages • media—Interface media changes • q921—Trace ISDN Q.921 frames • q931—Trace ISDN Q.931 frames 	<p>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"> • all—All interface tracing operations • event—Interface events • ipc—Interface interprocess communication (IPC) messages • media—Interface media changes • q921—Trace ISDN Q.921 frames • q931—Trace ISDN Q.931 frames 	<p>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • action-profile—Trace action profile invocation events. • all—Trace all events. • configuration—Trace configuration events. • protocol—Trace protocol processing events. • routing socket—Trace routing socket events. 	<p>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • all • change-events—Log changes that produce configuration events • config-states—Log the configuration state machine changes • kernel—Log configuration IPC messages to kernel • kernel-detail—Log details of configuration messages to kernel
match	—(Optional) Regular expression for lines to be traced.		—(Optional) Refine the output to log only those lines that match the given regular expression.	

Table 64: 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 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.</p>		<p>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option.</p> <p>Syntax: xk to specify KB, xm to specify MB, or xg to specify GB</p> <p>Range: 10 KB through 1 GB</p> <p>Default: 128 KB</p> <p>Default: If you do not include this option, tracing output is appended to an existing trace file.</p>	

Table 64: 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 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.</p> <p>Syntax: xk to specify kilobytes, xm to specify megabytes, or xg to specify gigabytes</p> <p>Range: 10 KB through the maximum file size</p>

Table 64: 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 <i>Default: 1 MB</i>
no-world-readable	—(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.
world-readable	—(Optional) Allow any user to read the log file.		—(Optional) Enable unrestricted file access.	—(Optional) Allow any user to read the log file.
disable				—(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 all .
no-remote-trace			—(Optional) Disable the remote trace.	-
match regex				—(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>
Tracing Operations of the Interface Process 397

reflective-relay

Syntax

```
reflective-relay;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family ethernet-switching]
```

Release Information

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D35 for the EX Series.

Description

Configure a switch interface to return packets back to a device on the same interface that was used to deliver the packets.

Default

Switch interfaces are not configured for reflective relay.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Reflective Relay for Use with VEPA Technology on QFX Switches

Configuring Reflective Relay on Switches

speed (Ethernet)

List of Syntax

[Syntax \(ACX Series, EX Series, MX Series\) on page 505](#)

[Syntax \(ACX5448, ACX710\) on page 505](#)

[Syntax \(ACX5448-D\) on page 505](#)

[Syntax \(EX Series\) on page 505](#)

[Syntax \(EX2300\) on page 505](#)

[Syntax \(EX4300\) on page 505](#)

[Syntax \(EX4600, OCX1100, QFX Series\) on page 505](#)

[Syntax \(QFX5100-48T\) on page 506](#)

Syntax (ACX Series, EX Series, MX Series)

```
speed (10m | 100m | 10g | 1g | 2.5g | 5g | auto | auto-10m-100m);
```

Syntax (ACX5448, ACX710)

```
speed (10g | 25g | 40g | 100g | 100m | 1g | auto);
```

Syntax (ACX5448-D)

```
speed (10g | 25g | 40g | 100g | 1g | auto);
```

Syntax (EX Series)

```
speed (auto-negotiation | speed) ;
```

Syntax (EX2300)

```
speed (10m | 100m | 1g | 2.5g );
```

Syntax (EX4300)

```
speed (10m | 100m | 1g | 2.5g | 5g | 10g);
```

Syntax (EX4600, OCX1100, QFX Series)

```
speed (10g | 1g | 100m)
```

Syntax (QFX5100-48T)

```
speed (10g | 1g | 100m | auto)
```

Hierarchy Level (ACX Series, EX Series, MX Series)

```
[edit interfaces interface-name],  
[edit interfaces ge-pim/0/0 switch-options switch-port port-number]
```

Hierarchy Level (EX Series)

```
[edit interfaces interface-name ether-options]
```

Hierarchy Level (ACX5448, ACX5448-D, ACX710, EX2300, EX4300, EX4600, OCX Series, QFX Series, QFX5100-48T)

```
[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 following interfaces:

- Management Ethernet interface (**fxp0** or **em0**)
- Fast Ethernet 12-port and 48-port PICs
- Built-in Fast Ethernet port on the FIC (M7i router)
- Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers
- Gigabit Ethernet interfaces on EX Series switches

If you enable autonegotiation, then the device automatically negotiates the speed based on the speed of the other end of the link. [Table 65 on page 507](#) describes the autonegotiation option available for different platforms:

Table 65: Autonegotiation Options

Autonegotiation Option	Description
auto-negotiation	Automatically negotiate the speed based on the speed of the other end of the link.
auto	Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link.
auto-10m-100m	Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link.

See [“Speed and Autonegotiation” on page 102](#) for more details.

Options

Routers support autonegotiation by default. To enable or disable autonegotiation, see *autonegotiation (Routers)*. [Table 66 on page 508](#) summarizes the speed and autonegotiation supported on different routing platforms:

Table 66: Speed and Autonegotiation support on Routers

Router	Speed	Autonegotiation Supported (Yes/No) ?
ACX Series	<ul style="list-style-type: none"> • 10M • 100M • 1G • 2.5G • 5G • 10G • auto • auto-10m-100m 	Yes
ACX5448, ACX710	<ul style="list-style-type: none"> • 10G • 25G • 40G • 100G • 100M • 1G • auto 	Yes
ACX5448-D	<ul style="list-style-type: none"> • 10G • 25G • 40G • 100G • 1G • auto 	Yes
MX Series	<ul style="list-style-type: none"> • 10M • 100M • 1G • 2.5G • 5G • 10G • auto • auto-10m-100m 	Yes

Switches support autonegotiation by default. To enable or disable autonegotiation, see [auto-negotiation \(Switches\)](#). [Table 13 on page 104](#) summarizes the speed and autonegotiation supported on different switching platforms:

Table 67: Speed and Autonegotiation support on Switches

Switch	Speed	Autonegotiation Supported (Yes/No) ?
EX Series	<ul style="list-style-type: none"> • 10M • 100M • 1G • 2.5G • 5G • 10G • auto • auto-10m-100m 	Yes
EX2300-48MP and EX2300-24MP	<ul style="list-style-type: none"> • 10M — supported on EX series switches and only on ge interfaces of EX2300MP switch. • 100M • 1G • 2.5G — supported only on mge interfaces of EX2300MP switch. 	Yes
EX4300-48MP (EX-UM-4SFPP-MR)	<ul style="list-style-type: none"> • 10M — supported only on ge interfaces. • 100M — supported on ge and mge interfaces. • 1G — supported on ge, mge 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. • 2.5G — supported only on mge interfaces. • 5G — supported only on mge interfaces. • 10G — supported on mge interfaces and 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet 	Yes

Table 67: Speed and Autonegotiation support on Switches (*continued*)

Switch	Speed	Autonegotiation Supported (Yes/No) ?
	uplink module on EX4300-48MP switches.	
EX4600, OCX1100, QFX Series, QFabric	<ul style="list-style-type: none"> • 100M • 1G • 10G 	No
QFX5100-48T	<ul style="list-style-type: none"> • 100M • 1G • 10G • auto <p>For more information, see “Configure Speed and Autonegotiation on QFX5100-48T Switches” on page 105.</p>	Yes
QFX5110-48S (connected with QFX-SFP-1GE-T transceiver)	<ul style="list-style-type: none"> • 100M • 1G • 10G • 40G • 100G • auto 	Yes
QFX5110-48S and QFX5100-48S (connected with JNP-SFPP-10GE-T transceiver)	<ul style="list-style-type: none"> • 100M • 1G • 10G <p>For more information, see “Autonegotiation support for QFX5110-48S and QFX5100-48S with JNP-SFPP-10GE-T transceiver” on page 106.</p>	Yes
QFX5120-48Y (connected with JNP-SFPP-10GE-T transceiver)	<ul style="list-style-type: none"> • 1G • 10G <p>For more information, see “Autonegotiation support for QFX5120-48Y with JNP-SFPP-10GE-T transceiver” on page 108.</p>	No

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Speed and Autonegotiation 102
Port Settings 130
Configuring the Interface Speed
Configuring the Interface Speed on Ethernet Interfaces
Configuring Gigabit Ethernet Autonegotiation
Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support 114
auto-negotiation 413
Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches 109
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

traps

Syntax

```
(traps | no-traps);
```

Hierarchy Level (ACX Series, MX Series, T Series, M Series, SRX Series, EX Series)

```
[edit dynamic-profiles profile-name interfaces interface-name],
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range name],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Hierarchy Level (QFX Series, EX4600)

```
[edit interfaces interface-name],
[edit interfaces interface-name unit logical-unit-number],
[edit interfaces interface-range interface-range-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.

Support at the **[edit dynamic-profiles *profile-name* interfaces *interface-name*]** hierarchy level introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description

Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

(Enhanced subscriber management for MX Series routers) To enable SNMP notifications, you must first configure the **interface-mib** statement at the **[edit dynamic-profiles *profile-name* interfaces *interface-name*]** hierarchy level. If **interface-mib** is not configured, the **traps** statement has no effect.

BEST PRACTICE: To achieve maximum performance when enhanced subscriber management is enabled, we recommend that you *not* enable SNMP notifications on all dynamic subscriber interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Enabling or Disabling SNMP Notifications on Physical Interfaces](#) | 63

[Enabling or Disabling SNMP Notifications on Logical Interfaces](#) | 82

unidirectional

Syntax

```
unidirectional;
```

Hierarchy Level

```
[edit interfaces interface-name],  
[edit logical-systems logical-system-name interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Create two new, unidirectional (transmit-only and receive-only) physical interfaces subordinate to the original parent interface. Unidirectional links are currently supported only on 10-Gigabit Ethernet interfaces on the following hardware:

- 4-port 10-Gigabit Ethernet DPC on the MX960 router
- 10-Gigabit Ethernet IQ2 PIC and 10-Gigabit Ethernet IQ2E PIC on the T Series router

Default

Disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Unidirectional Traffic Flow on Physical Interfaces

Enabling Unidirectional Traffic Flow on Physical Interfaces

unit

Syntax

The **ethernet-switching** and **fibre-channel** statements and all of their substatements are not supported on OCX Series switches.

```
unit logical-unit-number {
  family {
    ethernet-switching {
      filter input filter-name;
      filter output filter-name;
      native-vlan-id vlan-id;
      port-mode mode;
      vlan {
        members [ (all | names | vlan-ids) ];
      }
    }
    fibre-channel {
      port-mode (f-port | np-port);
    }
    inet {
      address address {
        primary;
      }
      filter input filter-name;
      filter output filter-name;
      primary;
      targeted-broadcast;
    }
  }
}
```

Hierarchy Level

```
[edit interfaces \(QFX Series\) interface-name],
[edit interfaces \(QFX Series\) interface-range interface-range-name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

NOTE: The **ethernet-switching** and **fibre-channel** statements and all of their substatements are not supported on OCX Series switches.

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Default

You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately. 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 and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 109](#)

[Configuring Link Aggregation | 281](#)

Junos OS Network Interfaces Library for Routing Devices

Configuration Statements: Gigabit Ethernet Interfaces

IN THIS CHAPTER

- [container-devices | 519](#)
- [craft-lockout | 520](#)
- [no-neighbor-learn | 522](#)

container-devices

Syntax

```
container-devices {  
    device-count number;  
}
```

Hierarchy Level

```
[edit chassis]  
[edit chassis interconnect-device name]  
[edit chassis node-group name]
```

Release Information

Statement introduced in Junos OS Release 11.3 for QFX Series switches.

Description

Specify the container devices configuration. The **number** option specifies the number of sequentially numbered container interfaces, from **ci0** to **ci127** maximum.

Options

number—Number of container devices.

Range: 1 through 128

Required Privilege Level

chassis—To view this statement in the configuration.

chassis-control—To add this statement to the configuration.

craft-lockout

Syntax

```
craft-lockout {
  alarm {
    interface-type {
      link-down (red | yellow | ignore);
    }
  }
  container-devices {
    device-count number;
  }
  fpc slot {
    pic pic-number {
      fibre-channel {
        port-range {
          port-range-low port-range-high;
        }
      }
    }
  }
  routing-engine
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
```

Hierarchy Level

```
[edit chassis -interconnect-device]
```

Release Information

Statement introduced in Junos Release 11.3 for the QFX Series.

Description

Disable the physical operation of the craft interface front panel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring the Junos OS to Disable the Physical Operation of the Craft Interface*

no-neighbor-learn

Syntax

```
no-neighbor-learn {
  accounting {
    destination-class-usage;
    source-class-usage direction;
  }
  address address;
  dhcp{
    client-identifier (ascii ascii | hexadecimal hexadecimal);
    lease-time (seconds | infinite);
    retransmission-attempt (DHCP Client) number;
    retransmission-interval (DHCP Client) seconds;
    server-address ip-address;
    update-server;
    vendor-id vendor-id;
  }
  filter {
    input filter-name;
    output filter-name;
  }
  mtu bytes;
  no-redirects;
  primary;
  rpf-check;
  targeted-broadcast;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Disable neighbor address learning on this interface (for both control plane and transit traffic) for all addresses or for the named address.

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 Routed VLAN Interfaces on Switches (CLI Procedure)

Configure ARP Learning and Aging Options

Configuration Statements: OTN Interfaces

IN THIS CHAPTER

- alarm low-light-alarm | 525
- encoding | 526
- fec | 527
- laser-enable | 528
- line-loopback | 529
- link-event-rate | 530
- modulation-format | 531
- optics-options | 532
- prbs | 533
- preemptive-fast-reroute | 534
- signal-degrade | 535
- tca | 536
- trigger | 538
- tx-power | 540
- warning | 541
- wavelength | 542

alarm low-light-alarm

Syntax

```
alarm low-light-alarm {
    (link-down | syslog);
}
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX 10016 switches.

Statement introduced in Junos OS Release 18.2R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Statement introduced in Junos OS Release 18.2R1 for PTX10K-LC1105 on the PTX10008 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

100-Gigabit Ethernet OTN Options Configuration Overview

encoding

Syntax

```
encoding (differential | non-differential);
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX10016 switches.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Description

Specify the encoding mode.

Options

differential—Differential phase line encoding.

non-differential—Non-differential phase line encoding.

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

100-Gigabit Ethernet OTN Options Configuration Overview

fec

Syntax

```
fec (efec | gfec | gfec-sdfec | hgfec | sd-fec | ufec | none);
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX10016 switches.

Description

Enable forward error correction (FEC) mode.

Default

The default value is **gfec**.

Options

efec—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

gfec—Generic forward error correction (GFEC) mode is configured to detect and correct bit errors.

gfec-sdfec—GFEC and soft-decision forward error correction (SD-FEC) modes are configured to detect and correct bit errors.

hgfec—High gain forward error correction mode is configured to detect and correct bit errors.

sd-fec—Soft-decision forward error correction mode is configured to detect and correct bit errors.

none—FEC mode is not configured.

ufec—Ultra Forward Error Correction (UFEC) mode is configured to detect and correct bit errors. This mode is supported only on 10G ports and not supported on the 40G and 100G ports.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[10-Gigabit Ethernet OTN Options Configuration Overview](#)

[100-Gigabit Ethernet OTN Options Configuration Overview](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#)

[Understanding Pre-FEC BER Monitoring and BER Thresholds](#)

laser-enable

Syntax

```
(laser-enable | no-laser-enable);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX10016 switches.

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](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#)

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 17.2R1 for QFX10008 and QFX10016 switches.

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](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#)

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 9.4 for EX Series switches.

Description

Configure the number of link fault management (LFM) events per second.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring Ethernet OAM Link Fault Management*

modulation-format

Syntax

```
modulation-format (qpsk | 8qam | 16qam);
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX10016 switches.

Statement introduced in Junos OS Release 18.3R1 for PTX10K-LC1104 on the PTX10008 and PTX10016 routers.

Description

Specify the modulation format.

Options

qpsk—Quadrature Phase Shift Keying modulation format.

8qam—8 quadrature amplitude modulation format.

16qam—16 quadrature amplitude modulation.

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

100-Gigabit Ethernet OTN Options Configuration Overview

optics-options

Syntax

```
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;
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX10016 switches.

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.

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 DWDM Interface Wavelength Overview](#)

[100-Gigabit Ethernet OTN Options Configuration Overview](#)

prbs

Syntax

```
(prbs | no-prbs);
```

Hierarchy Level

```
[edit interfaces interface-name otn-options]
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX10016 switches.

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](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#)

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 17.2R1 for QFX10008 and QFX 10016 switches.

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](#)

[100-Gigabit Ethernet OTN Options Configuration Overview](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#)

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 17.2R1 for QFX10008 and QFX 10016 switches.

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](#)

[100-Gigabit Ethernet OTN Options Configuration Overview](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#)

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 17.2R1 for QFX10008 and QFX10016 switches.

Command introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 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:

- Laser power
- Laser temperature

Default

By default, TCAs are not enabled.

Options

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
- **rx-power-high-tca**—Rx power high threshold setting trigger
- **rx-power-low-tca**—Rx power low threshold setting trigger

- **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

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

[optics-options](#) | **634**

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 17.2R1 for QFX10008 and QFX 10016 switches.

Description

Specify defect triggers.

Default

By default, triggers are ignored.

Options

trigger-identifier—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-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.

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](#)

[100-Gigabit Ethernet OTN Options Configuration Overview](#)

[Configuring OTN Interfaces on P1-PTX-2-100G-WDM](#)

tx-power

Syntax

```
tx-power dbm;
```

Hierarchy Level

```
[edit interfaces interface-name optics-options]
```

Release Information

Statement introduced in Junos OS Release 17.2R1 for QFX10008 and QFX10016 switches.

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

[optics-options](#) | **634**

100-Gigabit Ethernet OTN Options Configuration Overview

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 17.2R1 for QFX10008 and QFX10016 switches.

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

[optics-options](#) | 634

100-Gigabit Ethernet OTN Options Configuration Overview

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

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

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

Configuration Statements: Aggregated Ethernet Interfaces

IN THIS CHAPTER

- 802.3ad | 549
- alarm (chassis) | 551
- aggregated-devices | 553
- aggregated-ether-options | 555
- backup-liveness-detection | 559
- backup-peer-ip | 560
- bfd-liveness-detection (LAG) | 561
- chassis (EX Series) | 564
- chassis (QFabric System) | 566
- chassis (QFX Series) | 569
- device-count | 571
- disable (Link Protection) | 572
- disk-failure-action | 573
- disable (Multicast Load Balancing) | 574
- dlb | 575
- ether-options | 577
- ether-type | 585
- ethernet | 586
- ethernet (Aggregated Devices) | 587
- fibre-channel (Alarm) | 588
- flow-control | 589
- force-up | 591
- hash-parameters | 592
- interconnect-device (Chassis) | 594
- iccp | 596
- lacp (802.3ad) | 598
- lacp (Aggregated Ethernet) | 600

- link-down | 603
- link-mode | 604
- link-protection | 606
- link-protection-sub-group (802.3ad) | 608
- link-protection-sub-group (aggregated-ether-options) | 609
- link-speed | 611
- liveness-detection | 614
- local-bias | 615
- local-ip-addr (ICCP) | 616
- local-minimum-links-threshold | 617
- management-ethernet (Alarm) | 619
- minimum-interval (Liveness Detection) | 620
- minimum-links | 621
- minimum-receive-interval (Liveness Detection) | 623
- multicast-loadbalance | 624
- multiservice | 626
- node-device (Chassis) | 628
- node-group (Chassis) | 629
- non-revertive (Chassis) | 631
- non-revertive (Interfaces) | 632
- on-disk-failure | 633
- optics-options | 634
- peer (ICCP) | 636
- periodic | 638
- port-priority | 640
- routing-engine | 641
- rx-buffers | 642
- session-establishment-hold-time | 644
- transmit-interval (Liveness Detection) | 645
- tx-buffers | 646

802.3ad

Syntax

Syntax (EX Series)

```
802.3ad {
  aex;
  (backup | primary);
  lacp {
    force-up;
    port-priority
  }
}
```

Syntax (NFX, OCX, and QFX Series)

```
802.3ad ae{x;
  lacp {
    force-up;
    (primary | backup);
  }
  port-priority;
}
```

Hierarchy Level (EX Series)

[edit **interfaces** *interface-name* **ether-options**]

Hierarchy Level (NFX, OCX, and QFX Series)

[edit **interfaces** *interface-name* **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

Specify the aggregated Ethernet logical interface number.

NOTE:

- The **port-priority** statement is not supported on QFabric systems.
- The **force-up** statement is not supported on QFX10002 switches.

Options

- **aex**—Name of the LAG. Aggregated Ethernet logical interface number.
- **backup**—Designate the interface as the backup interface for link-protection mode.
- **primary**—Designate the interface as the primary interface for link-protection mode.

The remaining statements are described 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

[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)

[Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 321](#)

[Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches | 365](#)

[Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

[Configuring Link Aggregation | 281](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)

[Troubleshooting an Aggregated Ethernet Interface](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

alarm (chassis)

Syntax

```
alarm {
  interface-type {
    alarm-name (ignore | red | yellow);
  }
}
```

Hierarchy Level

```
[edit chassis],
[edit chassis interconnect-device name],
[edit chassis node-group name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 12.2 for the ACX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the chassis alarms and whether they trigger a red or yellow alarm, or whether they are ignored. Red alarm conditions light the **RED ALARM** LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the contact on the craft interface or LCD screen. Yellow alarm conditions light the **YELLOW ALARM** LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the craft interface or LCD screen.

To configure more than one alarm, include multiple **alarm-name** lines.

Options

alarm-name—Alarm condition. For a list of conditions, see *Configurable PIC Alarm Conditions*.

ignore—The specified alarm condition does not set off any alarm.

interface-type—Type of interface on which you are configuring the alarm: **atm**, **ethernet**, **sonet**, or **t3**.

red—The specified alarm condition sets off a red alarm.

yellow—The specified alarm condition sets off a yellow alarm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Understand Alarms</i>
<i>Chassis Conditions That Trigger Alarms</i>
<i>Chassis Alarm Messages on a QFX3500 Device</i>
<i>Interface Alarm Messages</i>

aggregated-devices

List of Syntax

[Syntax \(EX Series\) on page 553](#)

[Syntax \(QFX Series, EX4600, OCX1100, NFX Series\) on page 553](#)

Syntax (EX Series)

```
aggregated-devices {
  ethernet (Aggregated Devices) {
    device-count number;
    lacp
  }
}
```

Syntax (QFX Series, EX4600, OCX1100, NFX Series)

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

Hierarchy Level (EX Series, QFX Series)

[edit [chassis \(EX Series\)](#)]

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.

Statement introduced in Junos OS Release 14.2R3

Description

Configure properties for aggregated devices on the switch.

The remaining statements are explained separately. See [CLI Explorer](#).

Default

Aggregated devices are disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)

[Configuring Link Aggregation | 281](#)

[Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch | 299](#)

Junos OS Network Interfaces Library for Routing Devices

[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)

Configuring Aggregated Ethernet Links (CLI Procedure)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

Junos OS Ethernet Interfaces Configuration Guide

aggregated-ether-options

List of Syntax

[Syntax \(EX, MX Series\) on page 555](#)

[Syntax \(NFX, QFX Series, EX4600, OCX1100, QFabric\) on page 555](#)

Syntax (EX, MX Series)

```
aggregated-ether-options {
  ethernet-switch-profile {
    tag-protocol-id;
  }
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    admin-key key;
    periodic interval;
    system-id mac-address;
  }
  (link-protection | no-link-protection);
  link-speed speed;
  local-bias;
  logical-interface-fpc-redundancy;
  (loopback | no-loopback);
  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);
  }
  minimum-links number;
  system-priority
}
```

Syntax (NFX, QFX Series, EX4600, OCX1100, QFabric)

The **fcoe-lag** and **mc-ae** statements are not supported on OCX Series switches.

```

aggregated-ether-options {
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  ethernet-switch-profile {
    tag-protocol-id;
    (fcoe-lag | no-fcoe-lag);
    (flow-control | no-flow-control);
    lacp mode {
      admin-key key;
      periodic interval;
      system-id mac-address;
      force-up;
    }
  }
  (link-protection | no-link-protection);
  link-speed speed;
  local-bias;
  local-minimum-links-threshold threshold-value;
  (loopback | no-loopback);
  mc-ae {
    chassis-id chassis-id;
    mc-ae-id mc-ae-id;
    mode (active-active);
    status-control (active | standby);
  }
  minimum-links number;
  rebalance-periodic;
  resilient-hash;
  source-address-filter filter;
  (source-filtering | no-source-filtering);
}

```

Hierarchy Level (EX Series, QFX Series)

[edit **interfaces** aex]

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 12.3R2.

Statements **fcoe-lag** and **no-fcoe-lag** introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Statements **force-up**, **lACP**, and **resilient-hash** introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Statement **local-minimum-links-threshold** introduced in Junos OS Release 14.1X53-D40 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.

NOTE:

- The **fcoe-lag** and **mc-ae** statements are not supported on OCX Series switches.
- The **force-up** statement is not supported on QFX10002 switches.
- The **resilient-hash** statement is not supported on QFX5200, QFX5210, or QFX10002 switches.

The remaining statements are explained separately. See [CLI Explorer](#).

Default

Options are not enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch | 328](#)

Junos OS Network Interfaces Library for Routing Devices

[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)

[Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 321](#)

Configuring Aggregated Ethernet Links (CLI Procedure)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

Junos OS Ethernet Interfaces Configuration Guide

backup-liveness-detection

Syntax

```
backup-liveness-detection {  
    backup-peer-ip ipv4-address;  
}
```

Hierarchy Level

```
[edit protocols iccp peer]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 13.2R1 for EX Series switches.

Statement introduced in Junos OS Release 15.2R1 for MX Series routers.

Description

Determine whether a peer is up or down by exchanging keepalive messages over the management link between the two Inter-Chassis Control Protocol (ICCP) peers.

When an ICCP connection is operationally down, the status of the peers hosting a multichassis link aggregation group (MC-LAG) is detected by sending liveness detection requests to each other. Peers must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, the liveness detection check fails, and a failure action is implemented. Backup liveness detection must be configured on both peers hosting the MC-LAG.

For more information on the ICCP failure scenarios and handling the failures, refer to *ICCP Failure Scenarios for EX9200 Switches*.

The remaining statement is explained separately. See [CLI Explorer](#).

NOTE: If backup liveness detection is configured, the peer status is always up when either the ICCP TCP Connection is established, or Bidirectional Forwarding Protocol (BFD) is configured and the peer is up. The backup liveness check is only performed when the ICCP connection is down.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring Multichassis Link Aggregation on MX Series Routers*

backup-peer-ip

Syntax

```
backup-peer-ip ipv4-address;
```

Hierarchy Level

```
[edit protocols iccp peer backup-liveness-detection]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 13.2R1 for EX Series switches.

Description

Specify the IP address of the peer being used as a backup peer in the Bidirectional Forwarding Detection (BFD) configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

bfd-liveness-detection (LAG)

Syntax

```

bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  holddown-interval milliseconds;
  local-address bfd-local-address;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  neighbor bfd-neighbor-address;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}

```

Hierarchy Level

[edit interfaces *aex* aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 13.3.

Description

Configure Bidirectional Forwarding Detection (BFD) timers and authentication for aggregated Ethernet interfaces.

Options

holddown-interval *milliseconds*— Specify a time limit, in milliseconds, indicating the time that a BFD session remains up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

Range: 0 through 255,000

Default: 0

local-address *bfd-local-address*— Specify the loopback address or the AE interface address of the source of the BFD session.

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.

minimum-interval *milliseconds*— Specify a minimum time interval after which the local routing device transmits a BFD packet and then expects to receive a reply from the BFD neighbor. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** statement.

Range: 1 through 255,000

minimum-receive-interval *milliseconds*— Specify the minimum time interval after which the routing device expects to receive a reply from the BFD neighbor.

Range: 1 through 255,000

multiplier *number*— Specify the number of BFD packets that were not received by the BFD neighbor before the originating interface is declared down.

Range: 1 through 255

neighbor *bfd-neighbor-address*— Specify the loopback address or the AE interface address of a remote destination to send BFD packets.

no-adaptation— Disable the BFD adaptation. Include this statement if you do not want the BFD sessions to adapt to changing network conditions. We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

version— Configure the BFD version to detect (BFD version 1) or autodetect (the BFD version).

NOTE: The **version** option is not supported on the QFX Series. Starting in Junos OS Release 17.2R1, a warning will appear if you attempt to use this command.

Default: automatic

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

<i>authentication</i>
<i>detection-time</i>
<i>transmit-interval</i>
Configuring Micro BFD Sessions for LAG 336
<i>Example: Configuring Independent Micro BFD Sessions for LAG</i>
<i>Understanding Independent Micro BFD Sessions for LAG</i>

chassis (EX Series)

Syntax

```
chassis {
    aggregated-devices {
        ethernet (Aggregated Devices) {
            device-count number;
        }
    }
    auto-image-upgrade;
    fpc slot {
        pic pic-number {
            sfplus {
                pic-mode mode;
            }
        }
        power-budget-priority priority;
    }
    lcd-menu {
        fpc slot-number {
            menu-item (menu-name | menu-option) {
                disable;
            }
        }
    }
    nssu {
        upgrade-group group-name {
            fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
            member (NSSU Upgrade Groups) member-id {
                fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
            }
        }
    }
    psu {
        redundancy {
            n-plus-n (Power Management);
        }
    }
    redundancy {
        graceful-switchover;
    }
}
```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure chassis-specific properties for the switch.

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

<i>Configuring Aggregated Ethernet Links (CLI Procedure)</i>
<i>Upgrading Software by Using Automatic Software Download for Switches</i>
<i>Configuring the LCD Panel on EX Series Switches (CLI Procedure)</i>
<i>Configuring Graceful Routing Engine Switchover in a Virtual Chassis</i>
<i>Configuring Power Supply Redundancy (CLI Procedure)</i>
<i>Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade</i>

chassis (QFabric System)

Syntax

```
chassis {
  interconnect-device {
    alarm {
      (ethernet (Alarm) | management-ethernet) {
        link-down (red | yellow | ignore);
      }
    }
  }
  container-devices {
    device-count number;
  }
  craft-lockout {
    alarm {
      interface-type {
        link-down (red | yellow | ignore);
      }
    }
    container-devices {
      device-count number;
    }
  }
  fpc slot {
    power (on | off);
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
  fpc slot {
    power (on | off);
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
node-group name {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
}
```

```

    }
  }
  alarm {
    interface-type {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count number;
  }
  node-device name {
    fibre-channel {
      port-range {
        port-range-low port-range-high;
      }
    }
    pic pic-number {
      fte {
        port port-number;
        port-range port-range-low port-range-high;
      }
      xe {
        port port-number;
        port-range port-range-low port-range-high;
      }
      xle {
        port port-number;
        port-range port-range-low port-range-high;
      }
    }
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}

```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure chassis-specific properties for the switch.

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.

chassis (QFX Series)

Syntax

```
chassis {
  routing-engine
  redundancy {
    failover {
      on-disk-failure {
        disk-failure-action (halt | reboot);
      }
      on-loss-of-keepalives;
    }
    graceful-switchover;
  }
  aggregated-devices {
    ethernet {
      device-count number;
    }
    alarm {
      interface-type {
        alarm-name (red | yellow | ignore);
      }
    }
  }
  forwarding-options profile-name {
    num-65-127-prefix value
  }
  fpc slot {
    auto-speed-detection disable
    pic pic-number {
      port port-number {
        tunnel-port port-number tunnel-services;
        channel-speed speed;
      }
      port-range port-range-low port-range-high {
        channel-speed speed;
      }
    }
  }
  maximum-ecmp next-hops;
}
```

Hierarchy Level

[edit]

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Statement introduced in Junos OS Release 14.2R3

Description

Configure chassis-specific properties for the switch.

The remaining statements are explained separately. 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

Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion

[Configuring Link Aggregation](#) | 281

device-count

Syntax (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System)

```
device-count number;
```

Hierarchy Level (EX Series)

[edit [chassis \(EX Series\) aggregated-devices ethernet \(Aggregated Devices\)](#)]

Hierarchy Level (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System)

[edit [chassis aggregated-devices ethernet](#)],
[edit [chassis node-group name aggregated-devices ethernet](#)]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Range updated in Junos OS Release 9.5 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 14.2R3

Description

Configure the number of aggregated Ethernet logical devices available to the switch.

Options

number—Maximum number of aggregated Ethernet logical interfaces on the switch.

Range: 1 through 32 for EX2200, EX3200, and standalone EX3300 switches and for EX3300 Virtual Chassis

Range: 1 through 64 for standalone EX4200, standalone EX4500, and EX6200 switches and for EX4200 and EX4500 Virtual Chassis

Range: 1 through 239 for EX8200 Virtual Chassis

Range: 1 through 255 for standalone EX8200 switches

Range: 1 through 480 for standalone EX9200 switches

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion

[Configuring Link Aggregation | 281](#)

[Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch | 299](#)

[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)

Configuring Aggregated Ethernet Links (CLI Procedure)

[Junos OS Network Interfaces Configuration Guide](#)

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

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

disk-failure-action

Syntax

```
disk-failure-action (halt | reboot);
```

Hierarchy Level

```
[edit chassis routing-engine on-disk-failure],  
[edit chassis node-group name routing-engine on-disk-failure],  
[edit chassis interconnect-device name routing-engine on-disk-failure]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Halt or reboot when the Routing Engine hard disk fails.

Options

halt—Stop the Routing Engine.

reboot—Reboot the Routing Engine.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Enabling a Routing Engine to Reboot on Hard Disk Errors*

disable (Multicast Load Balancing)

Syntax

```
disable;
```

Hierarchy Level

```
[edit chassis multicast-loadbalance]
```

Release Information

Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description

(EX8200 switches only) Disable multicast load balancing. After you configure this statement and commit it, multicast load balancing no longer balances traffic on aggregated 10-gigabit Ethernet interfaces that are currently configured as well as those that are configured later.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)

dlb

Syntax

```
dlb {
  assigned-flow;
  per-packet;
  flowlet inactivity-interval;
  ether-type (ipv4|ipv6|mpls);
}
```

Hierarchy Level

[edit interfaces aex aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 19.4R1 for QFX5120-32C and QFX5120-48Y switches.

Description

Configure dynamic load balancing (DLB) on LAG with specific mode type—assigned flow, per-packet, or flowlet) and ether type—ipv4, ipv6, or mpls.

NOTE: Since ether-type is a global configuration and applicable to all LAGs in the system with DLB enabled, the CLI for specifying ether-type for LAG is added in the global format under forwarding-options. See [ether-type](#) for details.

Options

assigned-flow—Fixed link assignment.

per-packet—Per-packet link assignment.

flowlet *inactivity-interval*—Minimum inactivity interval for link re-assignment.

Range: 16 through 65535 (in micro seconds).

ether-type (*ipv4* | *ipv6* | *mpls*)—EtherType for DLB:

- *ipv4*—Set dynamic load balancing for IPv4.
- *ipv6*—Set dynamic load balancing for IPv6.
- *mpls*—Set dynamic load balancing for MPLS.

Required Privilege Level

interface - To view this statement in the configuration.

interface-control - To add this statement to the configuration.

RELATED DOCUMENTATION

[Dynamic Load Balancing | 372](#)

[Configuring Dynamic Load Balancing | 374](#)

[Example: Configure Dynamic Load Balancing | 376](#)

ether-options

List of Syntax

[Junos OS Syntax on page 577](#)

[Junos OS Evolved Syntax on page 578](#)

Junos OS Syntax

```
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
      (primary | backup);
      port-priority
    }
  }
  asynchronous-notification;
  (auto-negotiation| no-auto-negotiation);
  autostate-exclude
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  ethernet-switch-profile {
    ethernet-policer-profile
    (mac-learn-enable | no-mac-learn-enable);
    recovery-timeout time-in-seconds;
    storm-control storm-control-profile;
    tag-protocol-id;
  }
  (flow-control | no-flow-control);
  ieee-802-3az-eee;
  ignore-l3-incompletes;
  link-mode (automatic | full-duplex | half-duplex);
  (loopback | no-loopback);
  mdi-mode (auto | force | mdi | mdix);
  mpls {
    pop-all-labels <required-depth (1 | 2 | all)>;
  }
  no-auto-mdix;
  redundant-parent (Interfaces) parent;
  source-address-filter name;
  (source-filtering| no-source-filtering);
```

```

speed {
    (auto-negotiation <auto-negotiate-10-100> | ethernet-100m | ethernet-10g | ethernet-10m | ethernet-1g);
}
}

```

Junos OS Evolved Syntax

```

ether-options {
    802.3ad {
        aex;
        (backup | primary);
        lacp {
            force-up;
            (primary | backup);
            port-priority
        }
    }
    asynchronous-notification;
    (auto-negotiation | no-auto-negotiation);
    autostate-exclude
    ethernet-switch-profile {
        ethernet-policer-profile
        (mac-learn-enable | no-mac-learn-enable);
        recovery-timeout time-in-seconds;
        storm-control storm-control-profile;
        tag-protocol-id;
    }
    fec (gigether)
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    loopback-remote;
    mpls {
        pop-all-labels <required-depth (1 | 2 | all)>;
    }
    source-address-filter name;
    (source-filtering | no-source-filtering);
}

```

Hierarchy Level

```

[edit interfaces interface-name]
[edit interfaces interface-range range]

```

Release Information

Statement introduced in Junos OS Release 9.0.

autostate-exclude option introduced in Junos OS Release 14.1x53-D40 for QFX5100 switches only.

fec and **loopback-remote** options introduced in Junos OS Evolved Release 20.1R1.

Description

Configure **ether-options** properties for a Gigabit Ethernet or 10-Gigabit Ethernet interface.

In Junos OS Evolved, when you configure **set interfaces *interface* ether-options 802.3ad *ae name*** at the same time as you apply a second configuration to the same interface at the **[edit interfaces *interface*]** hierarchy, the second configuration will not take effect until the interface joins the aggregated Ethernet interface ***ae name***.

NOTE: The **ether-options** statement is not supported for subscriber management on aggregated Ethernet member link interfaces. You must configure **gether-options** instead.

[Table 68 on page 581](#) shows the supported and unsupported platforms.

Table 68: Supported Platform Information

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
ACX Series Routers (Junos OS) <ul style="list-style-type: none"> • ACX500 • ACX1000 • ACX1100 • ACX2100 • ACX2200 • ACX4000 • ACX5400 • ACX6000 	ACX Series Routers (Junos OS) <ul style="list-style-type: none"> • ACX5048 • ACX5096 	None
EX Series Switches (Junos OS) <ul style="list-style-type: none"> • EX2300 • EX2300 Multigigabit • EX2300-C • EX3400 • EX4300 • EX4300 Multigigabit • EX4600 • EX4650 • EX9200 • EX9250 	EX Series Switches (Junos OS) <ul style="list-style-type: none"> • EX2300 • EX2300 Multigigabit • EX2300-C • EX3400 • EX4300 • EX4300 Multigigabit • EX4600 • EX4650 • EX9200 • EX9250 	EX Series Switches (Junos OS) EX Series switches support both ether-options and gigether-options . To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use ether-options . To configure Forward Error Correction (FEC), use gigether-options .
MX Series Routers (Junos OS) <ul style="list-style-type: none"> • MX5 • MX10 • MX40 • MX80 • MX104 • MX150 • MX204 • MX240 • MX480 • MX960 • MX20008 	Not Supported	None

Table 68: Supported Platform Information (*continued*)

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
<ul style="list-style-type: none"> • MX2010 • MX2020 • MX10003 • MX10008 and MX10016 		
PTX Series Routers (Junos OS) <ul style="list-style-type: none"> • PTX1000 • PTX3000 • PTX5000 • PTX10001 • PTX10002 • PTX10003 • PTX10008 and PTX10016 	PTX Series Routers (Junos OS) <ul style="list-style-type: none"> • PTX1000 • PTX10001 • PTX10002 • PTX10003 • PTX10008 and PTX10016 	PTX Series Routers (Junos OS) <p>In Junos OS Release 17.3R3S7, PTX1000 Series routers support both ether-options and gigether-options.</p> <p>In Junos OS Releases 17.3R1, 17.4R1, and 17.4R2, PTX10000 Series routers support both ether-options and gigether-options.</p>
PTX Series routers (Junos OS Evolved) <ul style="list-style-type: none"> • PTX10003 • PTX10008 and PTX100016 	PTX Series routers (Junos OS Evolved) <ul style="list-style-type: none"> • PTX10003 • PTX10008 and PTX100016 	PTX Series routers (Junos OS Evolved) <p>Starting in Junos OS Evolved Release 20.1R1, PTX Series routers support ether-options only.</p>
		QFX Series Switches (Junos OS) <p>QFX Series switches support both ether-options and gigether-options.</p> <p>To configure Ethernet configuration options such as loopback, flow-control, auto-negotiation etc., use ether-options.</p> <p>To configure Forward Error Correction (FEC), use gigether-options.</p>

Table 68: Supported Platform Information (*continued*)

Supported Platforms for gigether-options	Supported Platforms for ether-options	Notes
QFX Series Switches (Junos OS) <ul style="list-style-type: none"> • QFX5100 (48S) • QFX5100 (48T) • QFX5100 (24Q) • QFX5100 (96S) • QFX5110 (48S) • QFX5110 (32Q) • QFX5120 (48Y) • QFX5120 (48T) • QFX5120 (48ym) • QFX5120 (32C) • QFX5200 (48Y) • QFX5200 (32C) • QFX5210 (64C) • QFX5210 (64C-S) • QFX5220 (32CD) • QFX5220 (128C) • QFX100002 • QFX10008 and QFX10016 	QFX Series Switches (Junos OS) <ul style="list-style-type: none"> • QFX5100 (48S) • QFX5100 (48T) • QFX5100 (24Q) • QFX5100 (96S) • QFX5110 (48S) • QFX5110 (32Q) • QFX5120 (48Y) • QFX5120 (48T) • QFX5120 (48ym) • QFX5120 (32C) • QFX5200 (48Y) • QFX5200 (32C) • QFX5210 (64C) • QFX5210 (64C-S) • QFX5220 (32CD) • QFX5220 (128C) • QFX100002 • QFX10008 and QFX10016 	
Not Supported	QFX Series Switches (Junos OS Evolved) <ul style="list-style-type: none"> • QFX5200-32C-L • QFX5220-32CD • QFX5220-128C 	None
SRX Series (Junos OS) <ul style="list-style-type: none"> • SRX300 • SRX550 • SRX1500 • SRX4100 and SRX4200 • SRX4600 • SRX5400 • SRX5600 	SRX Series (Junos OS) <ul style="list-style-type: none"> • SRX300 • SRX550 • SRX1500 • SRX4100 and SRX4200 • SRX4600 • SRX5400 • SRX5600 	SRX Series (Junos OS) <p>To configure gigabit-Ethernet interfaces (ge-), use gigether-options. To configure ethernet interfaces (et-) and fast ethernet interfaces (fe-), use ether-options.</p>

Default

Enabled.

Options

NOTE: The **auto-negotiation** and **speed** statements are not supported on the OCX Series.

loopback-remote—Starting in Junos OS Evolved Release 20.1R1, enable remote loopback.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Gigabit Ethernet Interface | 102](#)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

gether-options

ether-type

Syntax

```
ether-type (ipv4|ipv6|mpls);
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key ecmp-dlb]
```

Release Information

Statement introduced in Junos OS Release 19.4R1 for QFX5120-32C and QFX5120-48Y switches.

Description

EtherType for dynamic load balancing.

Options

ether-type (ipv4 | ipv6 | mpls)—Ether-type for DLB:

- **ipv4**—Set dynamic load balancing for IPv4.
- **ipv6**—Set dynamic load balancing for IPv6.
- **mpls**—Set dynamic load balancing for MPLS.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[dlb | 575](#)

[Dynamic Load Balancing | 372](#)

[Configuring Dynamic Load Balancing | 374](#)

[Example: Configure Dynamic Load Balancing | 376](#)

ethernet

Syntax

```
ethernet {  
    device-count number;  
}
```

Hierarchy Level

```
[edit chassis aggregated-devices],  
[edit chassis node-group aggregated-devices]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Statement introduced in Junos OS Release 14.2R3

Description

Configure properties for aggregated Ethernet devices on the switch.

The remaining statement is explained separately. 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

Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion

[Configuring Link Aggregation | 281](#)

Junos OS Network Interfaces Library for Routing Devices

ethernet (Aggregated Devices)

Syntax

```
ethernet {  
  device-count number;  
  lacp {  
    link-protection {  
      non-revertive;  
    }  
    system-priority;  
  }  
}
```

Hierarchy Level

[edit [chassis \(EX Series\) aggregated-devices](#)]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure properties for Ethernet aggregated devices on the switch.

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 Aggregated Ethernet Links \(CLI Procedure\)](#)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches](#) | 313

[Junos OS Ethernet Interfaces Configuration Guide](#)

fibre-channel (Alarm)

Syntax

```
fibre-channel {  
    link-down (red | yellow | ignore);  
}
```

Hierarchy Level

```
[edit chassis alarm],  
[edit chassis interconnect-device name alarm],  
[edit chassis node-group name alarm]
```

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure alarms for a Fibre Channel interface.

Options

The remaining statement is explained separately.—

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

flow-control

Syntax

```
(flow-control | no-flow-control);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.

You can configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces *interface-name* ether-options** hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.

NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).

Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

OCX Series switches do not support PFC.

- **flow-control**—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.
- **no-flow-control**—Disable flow control.

Default

Flow control is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[configured-flow-control](#) | 420

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Junos OS Network Interfaces Library for Routing Devices

force-up

Syntax

```
force-up;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options 802.3ad lacp]
```

Release Information

Statement introduced in Junos OS Release 10.0 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 16.1 for EX9200 switches.

Description

Set the state of the interface as UP when the peer has 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

[Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch | 328](#)

Junos OS Network Interfaces Library for Routing Devices

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

[Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support | 114](#)

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

[Junos OS Ethernet Interfaces Configuration Guide](#)

hash-parameters

Syntax

```
hash-parameters {
  ecmp {
    function {
      (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
    }
    offset offset;
    preprocess;
  }
  lag {
    function {
      (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
    }
    offset offset;
    preprocess;
  }
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key],
[edit routing-instances name forwarding-options enhanced-hash-key],
```

Release Information

Statement introduced in Junos OS Release 19.1R1 for QFX Series switches.

Description

Set traffic hashing parameters for ECMP or LAG traffic.

Options

ecmp —Set hashing parameters for ECMP traffic.

function (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo)—Set the hash functions for ECMP traffic or LAG traffic.

Values:

- **crc16-bisync**—Use CRC16-BISYNC function for ECMP or LAG traffic received on xe.
- **crc16-ccitt**—Use the CRC16-CCITT function for ECMP or LAG traffic received on xe.
- **crc32-hi**—Use the CRC32-HI function for ECMP or LAG traffic received on xe.
- **crc32-lo**—Use the CRC32-LO function for ECMP or LAG traffic received on xe.

lag—Set hashing parameters for LAG traffic.

offset *offset*—Set the hashing offset for ECMP or LAG traffic.

Range: 0 through 199

preprocess—Enable or disable the preprocess parameter for ECMP or LAG traffic.

Required Privilege Level

interface

RELATED DOCUMENTATION

| [Aggregated Ethernet Interfaces](#) | 261

interconnect-device (Chassis)

Syntax

```

interconnect-device {
  alarm {
    (ethernet (Alarm) | management-ethernet) {
      link-down (red | yellow | ignore);
    }
  }
  container-devices {
    device-count number;
  }
  craft-lockout {
    alarm {
      interface-type {
        link-down (red | yellow | ignore);
      }
    }
    container-devices {
      device-count number;
    }
  }
  fpc slot {
    power (on | off);
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
  fpc slot {
    power (on | off);
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}

```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties specific to a QFabric system Interconnect device.

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

| *Understanding Interconnect Devices*

iccp

Syntax

```
iccp {
  authentication-key string;
  local-ip-addr local-ip-addr;
  peer ip-address{
    authentication-key string;
    backup-liveness-detection {
      backup-peer-ip ip-address;
    }
    liveness-detection {
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    local-ip-addr ipv4-address;
    session-establishment-hold-time seconds;
  }
  session-establishment-hold-time seconds;
  traceoptions {
    file <filename> <files number> <match regular-expression> <microsecond-stamp> <size size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure Inter-Chassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

NOTE: Backup liveness detection is not supported on MX Series routers.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

lacp (802.3ad)

Syntax

```
lacp {  
    force-up;  
    (primary | backup);  
    port-priority;  
}
```

Hierarchy Level (EX Series)

[edit **interfaces** *interface-name* **ether-options** 802.3ad]

[edit interfaces aeX aggregated-ether-options]

[edit chassis aggregated-devices ethernet]

Hierarchy Level (QFX Series)

[edit **interfaces** *interface-name* **ether-options** 802.3ad]

Release Information

Statement introduced in Junos OS Release 10.0 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Support for LACP link protection introduced in Junos OS Release 11.4 for EX Series switches.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the Link Aggregation Control Protocol (LACP) parameters for aggregated Ethernet interfaces on the global level (for all the aggregated Ethernet interfaces on the switch) or for a specific aggregated Ethernet interface.

Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.

NOTE: The port-priority statement is not supported on QFabric systems.

NOTE: The force-up statement is not supported on QFX10002 switches.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)

[Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 321](#)

[Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

[Junos OS Ethernet Interfaces Configuration Guide](#)

[Configuring Link Aggregation | 281](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)

lACP (Aggregated Ethernet)

List of Syntax

[Syntax \(NFX Series\) on page 600](#)

[Syntax \(EX Series\) on page 600](#)

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 | 281](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)

Configuring LACP for Aggregated Ethernet Interfaces

link-down

Syntax

```
link-down (red | yellow | ignore);
```

Hierarchy Level

```
[edit chassis alarm ethernet (Alarm)],  
[edit chassis alarm fibre-channel],  
[edit chassis interconnect-device name alarm ethernet (Alarm)],  
[edit chassis node-group name alarm fibre-channel]
```

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify either red, yellow, or ignore to display when the link is down.

Options

red—Indicates that one or more hardware components have failed or exceeded temperature thresholds, or an alarm condition configured on an interface has triggered a critical warning.

yellow—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

ignore—Suppresses or ignores the alarm.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

link-mode

Syntax

```
link-mode mode;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Set the device's link-connection characteristic.

NOTE: Some QFX Series switch interfaces are full-duplex only. In such cases, the CLI shows only "link-mode full-duplex" or does not include the "link-mode" statement at all. You cannot configure the link mode on these interfaces; it is set to the full-duplex default.

Default

The **full-duplex** mode is enabled.

Options

mode —Link characteristic:

- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.
- **automatic**—Link mode is negotiated.

If **no-auto-negotiation** is specified in the **ether-options** option, you can select only **full-duplex** or **half-duplex**.

If **auto-negotiation** is specified in the **ether-options** option, you can select any mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | 109

Junos OS Network Interfaces Library for Routing Devices

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 gigether-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

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches](#) | 313

link-protection-sub-group (802.3ad)

Syntax

```
link-protection-sub-group group-name;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options 802.3ad]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description

Add an interface in an aggregated Ethernet bundle into a link-protection subgroup.

A link protection subgroup is created and named using the **link-protection-sub-group** statement in the **[edit interfaces aex aggregated-ether-options]** hierarchy.

Options

group-name—Name of the link protection subgroup that will include this interface after this statement is entered. The link protection subgroup is named when it is created using the **link-protection-sub-group** statement in the **[edit interfaces aex aggregated-ether-options]** hierarchy.

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 Links \(CLI Procedure\)](#)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches | 313](#)

[Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection](#)

link-protection-sub-group (aggregated-ether-options)

Syntax

```
link-protection-sub-group group-name {
    [primary | backup];
}
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description

Create and name a link protection subgroup.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a link aggregation group (LAG). If you need to provide link protection to a single link in a LAG, you do not need to configure link protection subgroups.

A link protection subgroup includes multiple links within the LAG. If one link in the primary link protection subgroup fails, traffic is forwarded over the links in the backup link protection subgroup.

Links within the LAG are added to the link protection subgroup using the [link-protection-sub-group](#) statement in the `[edit interfaces interface-name ether-options 802.3ad]` hierarchy.

Options

group-name—User-provided name of the link protection subgroup.

primary—Subgroup is the primary subgroup.

backup—Subgroup is the backup subgroup.

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 Links (CLI Procedure)

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches](#) | 313

| *Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection*

link-speed

Syntax

```
link-speed speed;
```

Hierarchy Level (QFX, NFX, EX Series, QFabric System, OCX1100, EX4600)

```
[edit interfaces aex aggregated-ether-options]
```

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 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

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.

speed—For aggregated Ethernet links, you can specify the speed in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

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

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.

NOTE: OCX Series switches only support 10g and 40g interfaces. Mixed rate aggregated Ethernet interfaces are not support on the OCX Series.

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 | 281](#)

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 | 291](#)

liveness-detection

Syntax

```
liveness-detection {
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
```

Hierarchy Level

[edit protocols **iccp peer**]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Enable Bidirectional Forwarding Detection (BFD). BFD enables rapid detection of communication failures between peers.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

local-bias

Syntax

```
local-bias;
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Release Information

Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices.

Description

Enable local link bias for all links in the aggregated Ethernet interface.

Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.

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

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Local Link Bias | 276](#)

[Understanding Local Link Bias | 274](#)

local-ip-addr (ICCP)

Syntax

```
local-ip-addr local-ip-address;
```

Hierarchy Level

```
[edit protocols iccp],  
[edit protocols iccp peer peer-IP-address]
```

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the local IP address of the interchassis link (ICL) interface that Inter-Chassis Control Protocol (ICCP) uses to communicate to the peers that host a multichassis link aggregation group (MC-LAG).

Options

local-ip-address—Default local IP address to be used by all peers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

local-minimum-links-threshold

Syntax

```
local-minimum-links-threshold threshold-value
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D40 for QFX5100 switches.

Description

For an aggregated Ethernet interface (LAG bundle) with member links spanning multiple chassis (member switches) in a Virtual Chassis or Virtual Chassis Fabric (VCF), set a threshold for the percentage of member links local to any particular chassis that must be up for that chassis to continue to be active in the aggregated Ethernet bundle.

NOTE: This statement is available only on member switches in a non-mixed mode QFX5100 Virtual Chassis or VCF.

You configure a threshold for a particular aggregated Ethernet interface (aex). When set, the threshold applies locally to any chassis with links in the specified LAG bundle. The threshold value represents a percentage of active local member links out of the total number of local member links for the chassis. When one or more member links on a chassis go down, the system compares the percentage of local member links that are “up” to the threshold. If the percentage of local member links that are “up” is less than the threshold, any remaining active local links are forced down as well, to prevent forwarding traffic for the aggregated Ethernet interface through any member links on that chassis. If the percentage of “up” links is greater than or equal to the threshold, the status of the active links remains unchanged and they can forward LAG traffic.

For example, consider a case where the threshold is set to 52, and one local LAG member link goes down on one switch in a Virtual Chassis Fabric that has a total of four local LAG member links. In this case, 75 percent of the links are still up (greater than the threshold, 52 percent), so the remaining local member links stay up. However, if two local member links go down, only 50 percent of the links are up, so the local minimum links feature forces the remaining two active local member links down as well.

This feature also adjusts local member link status accordingly when failed links come up again, if you reconfigure the threshold value, or if you add or remove local member links in the LAG bundle.

The local minimum links feature helps avoid traffic loss due to asymmetric bandwidth on the forwarding paths across a chassis when some local aggregated Ethernet member links fail and some remain active. Enable this feature only if you want to closely manage ingress and egress traffic forwarding paths on aggregated Ethernet interfaces, especially where local link bias is also enabled.

Default

The local minimum links feature is disabled by default.

Options

threshold-value—Percentage of member links in an aggregated Ethernet bundle local to a chassis that must be up for any local member links on that chassis to be active in the aggregated Ethernet bundle.

Range: 1 through 100 (decimal)

Default: none—This option is not enabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Understanding Local Minimum Links](#) | 276

management-ethernet (Alarm)

Syntax

```
management-ethernet {  
    link-down (red | yellow | ignore);  
}
```

Hierarchy Level

```
[edit chassis alarm],  
[edit chassis interconnect-device name alarm],  
[edit chassis node-group name alarm]
```

Release Information

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure alarms for a management Ethernet interface.

NOTE: If you configure a yellow alarm on the Interconnect device, it will be handled as a red alarm.

Options

The remaining statement is explained separately.—

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

minimum-interval (Liveness Detection)

Syntax

```
minimum-interval milliseconds;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure simultaneously the minimum interval at which the peer transmits liveness detection requests and the minimum interval at which the peer expects to receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately by using the **transmit-interval**, **minimal-interval** and **minimum-receive-interval** statements, respectively.

Options

milliseconds—Specify the minimum interval value for Bidirectional Forwarding Detection (BFD).

Range: 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

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.

(T Series, MX Series, PTX Series routers) You cannot configure the minimum number of links and the minimum bandwidth of an aggregated Ethernet bundle at the same time. They are mutually exclusive. To determine the status of the bundle, the device compares the value configured for minimum links and the value for minimum bandwidth. Because both cannot be configured at the same time, the device compares the configured value of the parameter with the default value of the other parameter. The device picks the higher value of the two parameters to determine the status of the Ethernet bundle. [Table 69 on page 622](#) describes how the device determines the bundle status based on sample values assigned to both parameters

Table 69: Determination of Bundle Status based on Minimum links and Minimum bandwidth Parameters

Current Bandwidth	Minimum Links	Minimum Bandwidth	Bundle Status
100G (10x10G)	2 Bandwidth=20G (2x10G)	N/A Default (1 bps)	Up
50G (5x10G)	N/A Default (1 link) Bandwidth = 10G	20G	Up
50G (5x10G)	N/A Default (1 link) Bandwidth = 10G	60G	Down

Options

number—Number of links.

Range: On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up. On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up. On EX8200 switches, the range of valid values for minimum-links number is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled up. On EX4600, QFX Series and Q Fabric Systems, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

Default: 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Aggregated Ethernet Minimum Links](#)

[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 | 291](#)

[Junos OS Services Interfaces Library for Routing Devices](#)

[Configuring Link Aggregation | 281](#)

minimum-receive-interval (Liveness Detection)

Syntax

```
minimum-receive-interval milliseconds;
```

Hierarchy Level

```
[edit protocols iccp peer liveness-detection]
```

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the minimum interval at which the peer must receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session.

Options

milliseconds—Specify the minimum interval value.

Range: 1 through 255,000

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

multicast-loadbalance

Syntax

```
multicast-loadbalance {
  disable (Multicast Load Balancing);
  hash-mode mode;
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description

(EX8200 switches only) Enable multicast load balancing to evenly distribute the Layer 3 routed multicast traffic on aggregated 10-gigabit Ethernet interfaces. When multicast load balancing is enabled on the EX8200 switch, it takes effect on aggregated 10-gigabit Ethernet interfaces that are currently configured as well as those that are configured later.

Crc-sgip mode is the default. If the default mode is implemented and the Layer 3 routed multicast traffic on the aggregated interfaces is not evenly distributed, you can try the other hashing algorithms. Use the algorithm that best balances your Layer 3 routed multicast traffic.

NOTE: If you change the hashing algorithm when multicast load balancing is disabled, the new algorithm takes effect after you reenable multicast load balancing.

Default

Multicast load balancing is enabled.

Options

hash-mode *mode*—Specify one of these hashing algorithms:

balanced—Use a round-robin method to determine the aggregated Ethernet links with the least amount of traffic.

crc-gip —Use multicast packets' group IP address as a key to a CRC algorithm to get the hash index. The hashing algorithm used is CRC.

crc-sgip—Use multicast packets' source IP address and group IP address for hashing. The hashing algorithm used is CRC.

crc-sip—Use multicast packets' source IP address for hashing. The hashing algorithm used is CRC.

simple-gip—Use multicast packets' group IP address for hashing. The hashing algorithm used is XOR.

simple-sgip—Use multicast packets' GIP bits to calculate the hash index. The hashing algorithm used is XOR.

simple-sip—Use multicast packets' source IP address bits for hashing. The hashing algorithm used is XOR.

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

[Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches | 365](#)

Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)

multiservice

Syntax

```
multiservice {
  source-mac;
  destination-mac;
  payload {
    ip {
      layer3-only;
      layer-3 (source-ip-only | destination-ip-only);
      layer-4;
      inner-vlan-id;
      outer-vlan-id;
    }
  }
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number hash-key family]
```

Release Information

Statement introduced in Junos OS Release 15.1X53-D10 for the QFX10000 switches.

Description

(QFX10000 switches only) Configure data used in a hash key for the **multiservice** protocol family when configuring PIC-level hashing for load balancing on an 802.3ad Link Aggregation Group.

Options

destination-mac—Include destination MAC address in the hash key.

payload—Include payload data in the hash key. This option has the following suboptions:

- **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key.
- **layer-3**—Include Layer 3 IP information in the hash key.
- **layer-4**—Include Layer 4 IP information in the hash key.
- **outer-vlan-id**—Include outer VLAN ID information in the hash key.
- **inner-vlan-id**—Include inner VLAN ID information in the hash key.

source-mac—Include source MAC address in the hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

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

node-device (Chassis)

Syntax

```
node-device name {
  fibre-channel {
    port-range {
      port-range-low port-range-high;
    }
  }
  pic pic-number {
    fte {
      port port-number;
      port-range port-range-low port-range-high;
    }
    xe {
      port port-number;
      port-range port-range-low port-range-high;
    }
  }
}
```

Hierarchy Level

[edit chassis [node-group](#)]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties specific to a Node device in a QFabric system.

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](#) | 281

node-group (Chassis)

Syntax

```
node-group name {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
  alarm {
    interface-type {
      link-down (ignore | red | yellow);
    }
  }
  container-devices {
    device-count number;
  }
  node-device name {
    fibre-channel {
      port-range {
        port-range-low port-range-high;
      }
    }
  }
  pic pic-number {
    fte {
      port port-number;
      port-range port-range-low port-range-high;
    }
    xe {
      port port-number;
      port-range port-range-low port-range-high;
    }
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure properties specific to a Node group.

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](#) | 281

non-revertive (Chassis)

Syntax

```
non-revertive;
```

Hierarchy Level

```
[edit chassis aggregated-devices ethernet 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.

Description

Disable the ability to switch to a better priority link (if one is available) once a link is established as active and a collection or distribution is enabled.

BEST PRACTICE: (MX Series) By default, Link Aggregation Control Protocol link protection is revertive. This means that after the current link becomes active, the router switches to a higher-priority link if one becomes operational or is added to the aggregated Ethernet bundle. In a highly scaled configuration over aggregated Ethernet, we recommend that you prevent the router from performing such a switch by including the **non-revertive** statement. Failure to do so may result in some traffic loss if a MIC on which a member interface is located reboots. Using the **non-revertive** statement for this purpose is not effective if both the primary and secondary interfaces are on the MIC that reboots.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Junos OS for Supporting Aggregated Devices

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches](#) | 313

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](#) | 606

Configuring Aggregated Ethernet Link Protection

[Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches](#) | 313

on-disk-failure

Syntax

```
on-disk-failure {  
    disk-failure-action (halt | reboot);  
}
```

Hierarchy Level

```
[edit chassis routing-engine],  
[edit chassis node-group name routing-engine],  
[edit chassis interconnect-device name routing-engine]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Halt or reboot the switch if it detects hard disk errors on the Routing Engine.

Options

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

Enabling a Routing Engine to Reboot on Hard Disk Errors

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

100-Gigabit Ethernet OTN Options Configuration Overview

Supported Forward Error Correction Modes on ACX6360 Router

peer (ICCP)

Syntax

```
peer ip-address {
  authentication-key string;
  backup-liveness-detection {
    backup-peer-ip ip-address;
  }
  liveness-detection {
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  local-ip-addr ipv4-address;
  session-establishment-hold-time seconds;
}
```

Hierarchy Level

[edit protocols [iccp](#)]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the peers that host a multichassis link aggregation group (MC-LAG). You must configure Inter-Chassis Control Protocol (ICCP) for both peers that host the MC-LAG.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

periodic

List of Syntax

[Syntax \(EX Series\) on page 638](#)

[Syntax \(QFX Series\) on page 638](#)

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

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch | 291](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

[Understanding Aggregated Ethernet Interfaces and LACP for Switches | 262](#)

Junos OS Network Interfaces Library for Routing Devices

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 | 313](#)

[Configuring Aggregated Ethernet LACP \(CLI Procedure\) | 283](#)

routing-engine

Syntax

```
routing-engine {  
  on-disk-failure {  
    disk-failure-action (halt | reboot);  
  }  
}
```

Hierarchy Level

```
[edit chassis]  
[edit chassis interconnect-device name],  
[edit chassis node-group name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure a Routing Engine to halt or reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting or halting prevents this.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Enabling a Routing Engine to Reboot on Hard Disk Errors

High Availability User Guide

rx-buffers

Syntax

```
rx-buffers (on | off);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options configured-flow-control]
```

Release Information

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.

Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the **rx-buffers** statement with the **tx-buffers** statement to configure asymmetric Ethernet PAUSE on an interface. (Use the **flow-control** statement to enable symmetric PAUSE and the **no-flow-control** statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)

NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).

Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

Default

Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.

Options

on | off—Enable or disable an interface to generate and send Ethernet PAUSE messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

flow-control 589
tx-buffers 646
<i>Configuring CoS Asymmetric Ethernet PAUSE Flow Control</i>
<i>Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control</i>
<i>Understanding CoS Flow Control (Ethernet PAUSE and PFC)</i>

session-establishment-hold-time

Syntax

```
session-establishment-hold-time seconds;
```

Hierarchy Level

```
[edit protocols iccp],  
[edit protocols iccp peer]
```

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Specify the time during which an Inter-Chassis Control Protocol (ICCP) connection must be established after IP route reachability between MC-LAG peers is up. When an MC-LAG peer detects IP route reachability to the MC-LAG peer, it tries to connect to it during the session-establishment-hold-time.

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.

Options

seconds—Time (in seconds) within which a successful ICCP connection must be established.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

transmit-interval (Liveness Detection)

Syntax

```
transmit-interval {  
    minimum-interval milliseconds;  
    threshold milliseconds;  
}
```

Hierarchy Level

[edit protocols **iccp peer liveness-detection**]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The negotiated transmit interval for a peer is the interval between the sending of BFD liveness detection requests to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

tx-buffers

Syntax

```
tx-buffers (on | off);
```

Hierarchy Level

```
[edit interfaces interface-name ether-options configured-flow-control]
```

Release Information

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.

Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the **tx-buffers** statement with the **rx-buffers** statement to configure asymmetric Ethernet PAUSE on an interface. (Use the **flow-control** statement to enable symmetric PAUSE and the **no-flow-control** statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)

NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).

Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

Default

Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.

Options

on | off—Enable or disable an interface to respond to an Ethernet PAUSE message.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

flow-control 589
rx-buffers 642
<i>Configuring CoS Asymmetric Ethernet PAUSE Flow Control</i>
<i>Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control</i>
<i>Understanding CoS Flow Control (Ethernet PAUSE and PFC)</i>

Configuration Statements: Channelizing Interfaces

IN THIS CHAPTER

- `channel-speed` | 649
- `fpc` | 650
- `fte (Port)` | 652
- `number-of-sub-ports` | 653
- `pic` | 655
- `pic-mode` | 658
- `sfppplus` | 659
- `short-reach-mode` | 660
- `xe (Port)` | 661
- `xle (Port)` | 663

channel-speed

Syntax

```
channel-speed (10g | 25g | 50g | 100g | disable-auto-speed-detection) ;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high)]
```

Release Information

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Enable the specified port on the Physical Interface Card (PIC) to perform in the specified channel speed. Additionally, you can disable auto-speed detection.

Default

The default speed of the platform is the maximum speed supported on the platform. Refer [“Port Settings” on page 130](#) for the supported speed details.

Options

10g—Set the channel speed to 10g (10-Gigabit Ethernet).

25g—Set the channel speed to 25g (25-Gigabit Ethernet).

50g—Set the channel speed to 50g (50-Gigabit Ethernet).

100g—Set the channel speed to 100g (100-Gigabit Ethernet).

disable-auto-speed-detection—Disable auto-speed detection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Port Settings](#) | 130

fpc

Syntax

```
fpc slot {
  auto-speed-detection disable;
  pic pic-number{
    tunnel-port port-number tunnel-services;
    port port-number{
      channel-speed (speed|disable-auto-speed-detection) ;
    }
    port-range port-range-low port-range-high {
      channel-speed (speed|disable-auto-speed-detection);
    }
  }
}
```

Hierarchy Level

[edit chassis]

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.

For generic routing encapsulation (GRE) tunneling, use the **tunnel-port** statement to specify the port that you want to convert to a GRE tunnel port.

Options

slot—Number of the FPC slot. For QFX3500, QFX3600, QFX5200, and OCX Series devices, the slot number is always 0.

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

| *show chassis fpc*

fte (Port)

Syntax

```
fte {
  (port port-number | port-range port-range-low port-range-high);
}
```

Hierarchy Level

```
[edit chassis \(QFX Series\) fpc slot pic pic-number]
```

Release Information

Statement introduced in Junos OS Release 13.2X52-D10 for the QFX Series.

Description

Configure a specific port or a range of ports to operate as 40-Gigabit Ethernet data plane uplink (*fte*) type ports.



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

Options

port-number—Port number on which you want to configure the port type. Valid values are 2 through 7.

port-range-low—Lowest-numbered port in the range of ports. The lowest possible value is 2.

port-range-high—Highest-numbered port in the range of ports. The highest possible value is 7.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the QSFP+ Port Type on QFX3500 Standalone Switches | 204](#)

[Configuring the Port Type on QFX3600 Standalone Switches | 202](#)

[Configuring the QSFP+ Port Type on QFX5100 Devices | 207](#)

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]
[edit interfaces interface-name]
```

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.

Hierarchy introduced in Junos OS Evolved Release 20.1R2 for JNP10K-LC1201 line cards on PTX10008 routers.

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** is 2, then two IFDs are created, namely et-x/y/z:0 and et-x/y/z:1.

The default value of channelized 10-Gigabit Ethernet interface is 4. The number of sub-ports that can be configured are, 1, 2, 3, or 4. You must set the **number-of-sub-ports** to 4, to channelize 40-Gigabit Ethernet interface to four 10-Gigabit Ethernet interfaces.

The **number-of-sub-ports** configuration statement can be used with rate selectability configuration at both PIC level and port level. This configuration statement is effective only when the port speed is 10 Gbps.

NOTE: You can configure the *number-of-sub-ports* only for 10-Gbps speed. For other speeds, this configuration is not supported.

(MPC11E) To specify the number of interfaces to be created on a physical port.

(JNP10K-LC1201) Specifies the number of channelized interfaces that you can configure on a physical port with the specified speed.

Options

number-of-sub-ports *number-of-subports*—Specify the number of sub-ports per physical port. For PTX10003-80C and PTX10003-160C routers, the values are 1, 2, and 4. On MPC10E-15C-MRATE line card, the values are 1, 2, 3, and 4. On JNP10K-LC1201, the values are 0 through 7.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[speed \(Ethernet\) | 505](#)

[PTX10003 Router Rate-Selectability Overview](#)

[Configuring Rate Selectability by Using New Port Profile Configuration](#)

pic

List of Syntax

[Syntax \(EX4600, OCX1100, QFX Series, QFabric system with ELS\) on page 655](#)

[Syntax \(EX4600, QFX Series, QFabric system\) on page 655](#)

Syntax (EX4600, OCX1100, QFX Series, QFabric system with ELS)

```
pic pic-number{
  tunnel-port port-number tunnel-services;
  port port-number{
    channel-speed (speed|disable-auto-speed-detection) ;
  }
  port-range port-range-low port-range-high {
    channel-speed (speed|disable-auto-speed-detection) ;
  }
}
```

Syntax (EX4600, QFX Series, QFabric system)

```
pic pic-number {
  fibre-channel {
    port-range {
      port-range-low port-range-high;
    }
  }
  fte{
    port port-number;
    (port-range port-range-low | port-range-high);
  }
  xe {
    (port port-number | port-range port-range-low port-range-high);
  }
  xle {
    (port port-number | port-range port-range-low port-range-high);
  }
}
```

Hierarchy Level (EX4600, OCX1100, QFX Series, QFabric system with ELS)

[edit [chassis \(EX Series\) fpc slot](#)]

Hierarchy Level (EX4600, QFX Series, QFabric system)

```
[edit chassis (QFX Series) fpc slot]
```

```
[edit chassis (QFX Series) node-group name node-device name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Options **xe** and **xle** introduced in Junos OS 12.2X50-D20 for the QFX Series.

Option **channel-speed** introduced in Junos OS Release 13.2 for the QFX Series.

Description

(QFX3500, QFX3600, and QFX5100 standalone switches running Enhanced Layer 2 Software only)—Configure a specific port or a range of ports to operate as 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.

Enable the specified port on the Physical Interface Card (PIC) to perform in the specified operating mode.

Options

pic *pic-number*—(QFX3500 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **1** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.

(QFX3600 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.

port *physical-port-number*—Port number on which you want to configure the port type.

port-range *port-range-low*—Lowest-numbered port in the range of ports.

port-range *port-range-high*—Highest-numbered port in the range of ports.

channel-speed (*speed* | **disable-auto-speed-detection)**—Configure *10g* for 10-Gigabit Ethernet type ports, and configure **disable-auto-speed-detection** to disable auto-channelization.

NOTE: This statement is not supported on the OCX Series.

- On a QFX3500 switch, specify **0** if the port type is **fiber-channel**, and **2** if the port type is **xle**.
- On a QFX3600 switch, specify **0** if the port type is **xe**, and **1** if the port type is **xle**.
- On a QFX5100 switch, specify **0** if the port type is **xe**, and **1** if the port type is **xle** and **fte**.

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

Channelizing Interfaces on QFX3500, QFX3600, QFX5100, QFX10002, QFX10008, QFX10016, and EX4600 Switches 134
Configuring the QSFP+ Port Type on QFX3500 Standalone Switches 204
Configuring the Port Type on QFX3600 Standalone Switches 202
Configuring the QSFP+ Port Type on QFX5100 Devices 207

pic-mode

Syntax

```
pic-mode mode;
```

Hierarchy Level

```
[edit chassis \(EX Series\) fpc slot pic pic-number sfpplus ]
```

Release Information

Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description

Configure the operating mode for the specified port on the SFP+ uplink module on an EX3200 or EX4200 switch.

On a switch using the SFP+ or SFP+ MACSec uplink module, the **pic-mode** setting defines the speeds of the SFP+ interfaces. When the PIC mode is set to **10g** on the SFP+ or SFP+ MACSec uplink module, uplink ports 0 and 2 support MACSec at 10-Gbps speeds while ports 1 and 3 cannot be used to send any traffic. When the PIC mode is set to **1g**, all four SFP+ ports on the uplink module support MACSec at 1-Gbps speeds.

Options

mode—Operating mode of the SFP+ uplink module:

- **1G**—1-gigabit operating mode
- **10G**—10-gigabit operating mode

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module](#) | 67

sfpplus

Syntax

```
sfpplus {
  pic-mode mode;
}
```

Hierarchy Level

```
[edit chassis (EX Series) fpc slot pic pic-number]
```

Release Information

Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description

Configure the operating mode for the specified port on the SFP+ uplink module on the EX3200 or EX4200 switch.

The remaining statement is explained separately. See [CLI Explorer](#).

Default

By default, the SFP+ uplink module operates in the 10-gigabit mode and supports SFP+ transceivers.

NOTE: The SFP+ uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module](#) | 67

short-reach-mode

Syntax

```
short-reach-mode (enable | disable);
```

Hierarchy Level

```
[edit chassis fpc fpc-slot pic pic-slot],  
[edit chassis fpc fpc-slot pic pic-slot port-range port-range-low port-range-high]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series.

Description

Configure short-reach mode for individual as well as a range of copper-based 10-Gigabit Ethernet interfaces using short cable lengths (less than 10m) on the QFX5100-48T switch. Short-reach mode reduces power consumption up to 5W on these interfaces.

Default

This feature is disabled by default.

Options

The following options are available:

- **enable**
- **disable**

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Short Reach Mode on QFX5100-48T](#) | 65

xe (Port)

Syntax

```
xe {
  (port port-number | port-range port-range-low port-range-high);
}
```

Hierarchy Level

[edit [chassis \(QFX Series\)](#) fpc slot [pic](#) *pic-number*]

Release Information

Statement introduced in Junos OS Release 12.2X50-D20 for the QFX Series.

Description

(QFX3600 standalone switch only) Configure a specific port or a range of ports to operate as four 10-Gigabit Ethernet (xe) type ports.



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

NOTE: Port Q0 supports only three (not the typical four) 10-Gigabit Ethernet ports. Therefore, you can configure up to 63 (not 64) 10-Gigabit Ethernet ports on ports Q0 through Q15.

Options

port-number—Port number on which you want to configure the port type. Valid values are 0 through 15.

port-range-low—Lowest-numbered port in the range of ports. The lowest possible value is 0.

port-range-high—Highest-numbered port in the range of ports. The highest possible value is 15.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Port Type on QFX3600 Standalone Switches](#) | 202

xle (Port)

Syntax

```
xle {
  (port port-number | port-range port-range-low port-range-high);
}
```

Hierarchy Level

```
[edit chassis \(QFX Series\) fpc slot pic pic-number]
```

Release Information

Statement introduced in Junos OS Release 12.2X50-D20 for the QFX Series.

Description

(QFX3500 and QFX3600 standalone switches only) Configure a specific QSFP+ port or a range of QSFP+ ports to operate as 40-Gigabit Ethernet (*xle*) type ports.



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

Options

port-number—Port number on which you want to configure the port type. On a QFX3500 standalone switch, specify a value from 0 through 3. On a QFX3600 standalone switch, specify a value from 0 through 15.

port-range-low—Lowest-numbered port in the range of ports. The lowest possible value is 0.

port-range-high—Highest-numbered port in the range of ports. The highest possible value is 3 on QFX3500 standalone switches, and 15 on QFX3600 standalone switches.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the QSFP+ Port Type on QFX3500 Standalone Switches](#) | 204

Configuration Statements: Energy Efficient Interfaces

IN THIS CHAPTER

- [ieee-802-3az-eee](#) | 665

ieee-802-3az-eee

Syntax

```
ieee-802-3az-eee;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
```

Release Information

Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description

Configure Energy Efficient Ethernet (EEE) on an EEE-capable Base-T copper interface.

Default

EEE is disabled on EEE-capable interfaces.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Energy Efficient Ethernet on Interfaces](#) | 211

Configuration Statements: VLANs

IN THIS CHAPTER

- `ethernet-switch-profile` | 667
- `l2-domain-id-for-l3` | 669
- `layer2 (enhanced-hash-key)` | 670
- `layer3-domain-identifier` | 672
- `members` | 673
- `no-local-switching` | 675
- `port-mode` | 676
- `tag-protocol-id (TPIDs Expected to Be Sent or Received)` | 678
- `vlan-id` | 680
- `vlan-tagging` | 681

ethernet-switch-profile

Syntax

```

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;
      }
    }
    storm-control storm-control-profile;
    tag-protocol-id tpid;
  }
  mac-learn-enable;
}

```

Hierarchy Level

```

[edit interfaces interface-name gigether-options],
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name aggregated-ether-options],
[edit interfaces interface-name ether-options]

```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 13.2X50-D15 for the EX Series switches.

Description

NOTE: On QFX Series standalone switches, the **ethernet-policer-profile** CLI hierarchy and the **mac-learn-enable** statement are supported only on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router); 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, configure VLAN tag and MAC address accounting and filtering properties.

The remaining statements are explained separately. See [CLI Explorer](#).

NOTE: When you gather interfaces into a bridge domain, the **no-mac-learn-enable** statement at the **[edit interfaces interface-name gigether-options ethernet-switch-profile]** hierarchy level is not supported. You must use the **no-mac-learning** statement at the **[edit bridge-domains bridge-domain-name bridge-options interface interface-name]** hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

Default

If the **ethernet-switch-profile** statement is not configured, 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) behave like 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

Configuring Gigabit Ethernet Policers

Configuring Gigabit Ethernet Policers

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

I2-domain-id-for-l3

Syntax

```
I2-domain-id-for-l3 id;
```

Hierarchy Level

```
[edit routing-instances instance-name]
```

Release Information

Statement introduced in Junos OS Release 12.3R2.

Description

Specify a Layer 2 domain ID within a routing instance.

Options

id—Layer 2 identification number.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring a Layer 2 Virtual Switch on an EX Series Switch*

layer2 (enhanced-hash-key)

List of Syntax

[Syntax \(EX Series and QFX5100 Switch\) on page 670](#)

[Syntax \(QFX10002 Switch\) on page 670](#)

Syntax (EX Series and QFX5100 Switch)

```
layer2 {
  no-destination-mac-address;
  no-ether-type;
  no-source-mac-address;
  vlan-id;
}
```

Syntax (QFX10002 Switch)

```
layer2 {
  no-incoming-port;
  no-destination-mac-address;
  no-ether-type;
  no-source-mac-address;
  source-mac-address;
  vlan-id;
  no-vlan-id;
  inner-vlan-id;
}
```

Hierarchy Level

[edit forwarding-options [enhanced-hash-key](#)]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

Description

Select the fields in the Layer 2 header that are used by the hashing algorithm to make hashing decisions.

When traffic enters a link aggregation group (LAG) bundle, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links. The hashing algorithm always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When traffic is exiting a device that has enabled ECMP, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to forward traffic to the next hop device.

The hashing algorithm only inspects the fields in the Layer 2 header when the hash mode is set to Layer 2 header. You can set the hash mode to Layer 2 header using the **set forwarding-options enhanced-hash-key hash-mode layer2-header** statement.

Default

The hash mode of the hashing algorithm is set to Layer 2 payload, by default. When the hash mode is set to Layer 2 payload, the hashing algorithm does not use fields in the Layer 2 header to make hashing decisions.

The following fields are used by the hashing algorithm when the hash mode of the hashing algorithm is set to Layer 2 header, by default:

- Destination MAC address
- Ethertype
- Source MAC address

Options

no-destination-mac-address—Exclude the destination MAC address field from the hashing algorithm.

no-ether-type—Exclude the Ethertype field from the hashing algorithm.

no-source-mac-address—Exclude the source MAC address field from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) | 349](#)

[Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic | 342](#)

[enhanced-hash-key | 725](#)

[hash-mode | 733](#)

layer3-domain-identifier

Syntax

```
layer3-domain-identifier identifier;
```

Hierarchy Level

```
[edit routing-instances instance-name]
```

Release Information

Statement introduced in Junos OS Release 12.3R2.

Description

Specify a Layer 3 domain ID within a routing instance.

Options

id—Layer 3 identification number.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring a Layer 2 Virtual Switch on an EX Series Switch*

members

Syntax

```
members [(all | names | vlan-ids)];
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family ethernet-switching vlan]
```

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement updated with enhanced ? (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

Description

For trunk interfaces, configure the VLANs that can carry traffic.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after **vlan** or **vlangs** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command **set vlans id vlan-id ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum.

On an EX Series switch that runs Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style, the maximum number of VLAN members allowed on the switch is 8 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 8). If the switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (**eswd**) due to memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 24). If the configuration of one of these switches exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

Options

all—Specifies that this trunk interface is a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.

NOTE: Since VLAN members are limited, specifying **all** could cause the number of VLAN members to exceed the limit at some point.

names—Name of one or more VLANs. VLAN IDs are applied automatically in this case.

NOTE: **all** cannot be a VLAN name.

vlan-ids—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, **10-20** or **10-20 23 27-30**.

NOTE: Each configured VLAN must have a specified VLAN ID to successfully commit the configuration; otherwise, the configuration commit fails.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

show ethernet-switching interfaces

show ethernet-switching interface

show vlans

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

[Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support | 114](#)

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

Configuring VLANs for EX Series Switches

Configuring VLANs for EX Series Switches with ELS Support (CLI Procedure)

no-local-switching

Syntax

```
no-local-switching;
```

Hierarchy Level

```
[edit routing-instances instance-name]
```

Release Information

Statement introduced in Junos OS Release 12.3R2.

Description

Specify that access ports in this routing instance do not forward packets to each other.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring a Layer 2 Virtual Switch on an EX Series Switch*

port-mode

Syntax

```
port-mode (access | tagged-access | trunk);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family ethernet-switching]
```

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.

Description

NOTE: This statement does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see [interface-mode](#). For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

Configure whether an interface on the switch operates in access, tagged access, or trunk mode.

Default

All switch interfaces are in access mode.

Options

access—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.

tagged-access—Have the interface operate in tagged-access mode. In this mode, the interface can be in multiple VLANs. Tagged access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.

trunk—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.

NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command **set vlans id vlan-id ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum. To determine the maximum number of VLAN members allowed on a switch, multiply the VLAN maximum for the switch times 8 (vmember limit = vlan max * 8).

If a switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (**eswd**) due to memory allocation failure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Connecting an EX Series Access Switch to a Distribution Switch

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

Configuring VLANs for EX Series Switches

[Junos OS Ethernet Interfaces Configuration Guide](#)

tag-protocol-id (TPIDs Expected to Be Sent or Received)

Syntax

```
tag-protocol-id [tpids];
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options ethernet-switch-profile],  
[edit interfaces interface-name aggregated-ether-options ethernet-switch-profile],  
[edit interfaces interface-name aggregated-ether-options ethernet-switch-profile],  
[edit interfaces interface-name ether-options ethernet-switch-profile]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Statement introduced in Junos OS Release 14.1X53-D15 for the QFX Series.

Description

For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with 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), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the **tag-protocol-id** statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.

For 10-Gigabit Ethernet LAN/WAN PIC interfaces on T Series routers only the default TPID value (**0x8100**) is supported.

For Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, define the TPIDs expected to be sent or received on a particular VLAN. The default TPID value is **0x8100**. Other supported values are **0x88a8**, **0x9100**, and **0x9200**.

Options

tpids—TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames

Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support

vlan-id

Syntax

```
vlan-id vlan-id-number;
```

Hierarchy Level

```
[edit interfaces \(QFX Series\) interface-name unit logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 9.2 for EX Series switches.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For 10-Gigabit Ethernet and aggregated Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface. Statement introduced in Junos OS Release 9.2 for EX Series switches.

NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

Options

vlan-id-number—Valid VLAN identifier.

Range: 1 through 4094

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[vlan-tagging](#) | **681**

[Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches](#) | **109**

[Configuring a Layer 3 Logical Interface](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

vlan-tagging

Syntax

```
vlan-tagging;
```

Syntax (QFX Series, NFX Series, and EX4600)

```
vlan-tagging;
```

Syntax (SRX Series Interfaces)

```
vlan-tagging native-vlan-id vlan-id;
```

Hierarchy Level

```
[edit interfaces interface-name],  
[edit logical-systems logical-system-name interfaces interface-name]
```

QFX Series, NFX Series, and EX4600 Interfaces

```
[edit interfaces \(QFX Series\) interface-name ]  
[edit interfaces \(QFX Series\) interface-range interface-range-name ]
```

SRX Series Interfaces

```
[edit interfaces interface ]
```

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 9.5.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement introduced in Junos OS Release 13.2 for PTX Series Routers.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description

For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

NOTE: For QFX Series configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface. Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.

On EX Series switches except for EX4300 and EX9200 switches, the **vlan-tagging** and **family ethernet-switching** statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default **family** setting.

Default

VLAN tagging is disabled by default.

Options

native-vlan-id— (SRX Series) Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.

NOTE: The **native-vlan-id** can be configured only when either **flexible-vlan-tagging** mode or **interface-mode** trunk is configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

802.1Q VLANs Overview

Configuring a Layer 3 Subinterface (CLI Procedure)

Configuring Tagged Aggregated Ethernet Interfaces

Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch

[vlan-id | 680](#)

Configuring a Layer 3 Logical Interface

Configuring VLAN Tagging

Configuration Statements: Uplink Failure Detection for Interfaces

IN THIS CHAPTER

- `debounce-interval` | 685
- `group` | 686
- `link-to-disable` | 687
- `link-to-monitor` | 688
- `uplink-failure-detection` | 689

debounce-interval

Syntax

```
debounce-interval seconds;
```

Hierarchy Level

```
[edit logical-systems name protocols uplink-failure-detection group],  
[edit protocols uplink-failure-detection group]
```

Release Information

Statement introduced in Junos OS Release 19.2R1 for the QFX Series.

Description

The amount of time, in seconds, that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces.

Options

seconds—Delay, in seconds, before bringing up the **link-to-disable** port.

Range: 5 to 300 seconds

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

group

Syntax

```
group group-name {  
    debounce-interval seconds;  
    link-to-monitor interface-name;  
    link-to-disable interface-name;  
}
```

Hierarchy Level

```
[edit protocols uplink-failure-detection]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the EX Series.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure a group of uplink and downlink interfaces for uplink failure detection.

Options

group-name—Name of the uplink failure detection group.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Uplink Failure Detection | 215](#)

[Configuring Interfaces for Uplink Failure Detection | 217](#)

[Example: Configuring Interfaces for Uplink Failure Detection | 218](#)

link-to-disable

Syntax

```
link-to-disable interface-name;
```

Hierarchy Level

```
[edit protocols uplink-failure-detection group group-name]
```

Release Information

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of eight downlink interfaces in a group.

Options

interface-name—Name of the downlink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Uplink Failure Detection | 215](#)

[Configuring Interfaces for Uplink Failure Detection | 217](#)

[Example: Configuring Interfaces for Uplink Failure Detection | 218](#)

link-to-monitor

Syntax

```
link-to-monitor interface-name;
```

Hierarchy Level

```
[edit protocols uplink-failure-detection group group-name]
```

Release Information

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of eight uplink interfaces in a group.

Options

interface-name—Name of the uplink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Uplink Failure Detection | 215](#)

[Configuring Interfaces for Uplink Failure Detection | 217](#)

[Example: Configuring Interfaces for Uplink Failure Detection | 218](#)

uplink-failure-detection

Syntax

```
uplink-failure-detection {  
  group group-name {  
    debounce-interval seconds;  
    link-to-monitor interface-name;  
    link-to-disable interface-name;  
  }  
}
```

Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description

Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failure information to the downlink interfaces.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Overview of Uplink Failure Detection | 215](#)

[Configuring Interfaces for Uplink Failure Detection | 217](#)

[Example: Configuring Interfaces for Uplink Failure Detection | 218](#)

Configuration Statements: Unicast Reverse Path Forwarding (uRPF)

IN THIS CHAPTER

- `group` (RPF Selection) | 691
- `multicast-rpf-routes` | 692
- `next-hop` (PIM RPF Selection) | 693
- `prefix-list` (PIM RPF Selection) | 694
- `rpf-check` (Dynamic Profiles) | 695
- `rpf-check` | 696
- `rpf-check-policy` (Routing Options RPF) | 698
- `rpf-loose-mode-discard` | 699
- `rpf-selection` | 700
- `source` (PIM RPF Selection) | 702
- `unicast-reverse-path` | 703
- `wildcard-source` (PIM RPF Selection) | 704

group (RPF Selection)

Syntax

```
group group-address{
  source source-address{
    next-hop next-hop-address;
  }
  wildcard-source {
    next-hop next-hop-address;
  }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name edit protocols pim rpf-selection]
```

Release Information

Statement introduced in JUNOS Release 10.4.

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

Configure the PIM group address for which you configure RPF selection [group \(RPF Selection\)](#).

Default

By default, PIM RPF selection is not configured.

Options

group-address—PIM group address for which you configure RPF selection.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

multicast-rpf-routes

Syntax

```
multicast-rpf-routes;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis traffic-engineering family inet shortcuts],  
[edit logical-systems logical-system-name routing-instances traffic-engineering family inet shortcuts],  
[edit protocols isis traffic-engineering family inet shortcuts],  
[edit routing-instances routing-instance-name protocols isis traffic-engineering family inet shortcuts]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Install unicast IPv4 routes into the multicast routing table (inet.2) for multicast reverse-path-forwarding (RPF) checks.

Traffic engineering shortcuts must be enabled. IPv4 multicast topology must not be enabled. Label-switched paths (LSPs) must not be advertised into IS-IS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Enabling IS-IS Traffic Engineering Support

Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts

next-hop (PIM RPF Selection)

Syntax

```
next-hop next-hop-address;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address source source-address],  
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address wildcard-source],  
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses source  
  source-address],  
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses  
  wildcard-source]
```

Release Information

Statement introduced in JUNOS Release 10.4.

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

Configure the specific next-hop address for the PIM group source.

Options

next-hop-address—Specific next-hop address for the PIM group source.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Example: Configuring PIM RPF Selection*

prefix-list (PIM RPF Selection)

Syntax

```
prefix-list prefix-list-addresses {
  source source-address {
    next-hop next-hop-address;
  }
  wildcard-source {
    next-hop next-hop-address;
  }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address source source-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address wildcard-source],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses source
source-address],
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses
wildcard-source]
```

Release Information

Statement introduced in Junos OS Release 10.4.

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

(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.

Options

prefix-list-addresses—List of prefixes (addresses) for multiple PIM groups.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

rpf-check (Dynamic Profiles)

Syntax

```
rpf-check {
    fail-filter filter-name;
    mode loose;
}
```

Hierarchy Level

```
[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Description

Reduce forwarding of IP packets that might be spoofing and address by checking whether traffic is arriving on an expected path that the sender would use to reach the destination. You can include this statement with the **inet** protocol family only. When the traffic passes the check, it is forwarded to the destination address; otherwise it is discarded. When you configure **rpf-check** alone, then unicast RPF is in strict mode, meaning that the check passes only when the packet's source address is in the FIB and the interface matches the routes RPF.

Starting in Junos OS Release 19.1, the **show interfaces statistics *logical-interface-name* detail** command displays unicast RPF statistics for dynamic logical interfaces when either **rpf-check** or **rpf-check mode loose** is enabled on the interface. No additional statistics are displayed when **rpf-check fail-filter *filter-name*** is configured on the interface. The **clear interfaces statistics *logical-interface-name*** command clears RPF statistics.

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

[Unicast RPF in Dynamic Profiles for Subscriber Interfaces](#)

[Configuring Unicast RPF Strict Mode](#)

rpf-check

List of Syntax

[Syntax \(MX Series, SRX Series, M Series, T Series, PTX Series\) on page 696](#)

[Syntax \(EX Series and QFX Series\) on page 696](#)

Syntax (MX Series, SRX Series, M Series, T Series, PTX Series)

```
rpf-check {
    fail-filter filter-name;
    mode loose;
}
```

Syntax (EX Series and QFX Series)

```
rpf-check;
```

Hierarchy Level (MX Series, SRX Series, M Series, T Series, PTX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit interfaces interface-name unit logical-unit-number family inet6],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6]
```

Hierarchy Level (EX Series and QFX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Support for interface **ps0** (pseudowire subscriber logical interface device) added in Junos OS Release 15.1.

Description

Enable a reverse-path forwarding (RPF) check on unicast traffic.

On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.

On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.

On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.

On QFX Series switches, enable an RPF check on unicast traffic on the selected ingress interfaces. ECMP packets are checked by QFX5000 Series switches only.

The mode statement is explained separately.

Default

Unicast RPF is disabled on all interfaces.

Options

fail-filter—A filter to evaluate when packets are received on the interface. If the RPF check fails, this optional filter is evaluated. If the fail filter is not configured, the default action is to silently discard the packet.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Unicast RPF Strict Mode

Configuring Unicast RPF Loose Mode

Configuring a Pseudowire Subscriber Logical Interface Device

Example: Configuring Unicast RPF (On a Switch)

rpf-check-policy (Routing Options RPF)

Syntax

```
rpf-check-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast],  
[edit logical-systems logical-system-name routing-options multicast],  
[edit routing-instances routing-instance-name routing-options multicast],  
[edit routing-options multicast]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.

Options

policy-names—Name of one or more multicast RPF check policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Example: Configuring RPF Policies*

rpf-loose-mode-discard

Syntax

```
rpf-loose-mode-discard {  
  family {  
    inet;  
    inet6;  
  }  
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Release Information

Statement introduced in Junos OS Release 12.1.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard next hop.

Options

inet—IPv4 address family.

inet6—IPv6 address family.

Required Privilege Level

interface-control—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Unicast RPF (Routers)

rpf-selection

Syntax

```
rpf-selection {  
  group group-address {  
    source source-address {  
      next-hop next-hop-address;  
    }  
  }  
  wildcard-source {  
    next-hop next-hop-address;  
  }  
}  
prefix-list prefix-list-addresses {  
  source source-address {  
    next-hop next-hop-address;  
  }  
  wildcard-source {  
    next-hop next-hop-address;  
  }  
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim]  
[edit protocols pim]
```

Release Information

Statement introduced in JUNOS Release 10.4.

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

Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.

NOTE: Starting in Junos OS 17.4R1, you can configure **rpf-selection** statement at the **[edit protocols pim]** hierarchy level.

The remaining statements are explained separately. See [CLI Explorer](#).

Default

If you omit the **rpf-selection** statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.

Options

source-address—Specific source address for the PIM group.

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Example: Configuring PIM RPF Selection*

source (PIM RPF Selection)

Syntax

```
source source-address {  
    next-hop next-hop-address;  
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address],  
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses]
```

Release Information

Statement introduced in JUNOS Release 10.4.

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

Configure the source address for the PIM group.

Options

source-address—Specific source address for the PIM group.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Example: Configuring PIM RPF Selection*

unicast-reverse-path

Syntax

```
unicast-reverse-path (active-paths | feasible-paths);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options forwarding-table],  
[edit routing-instances routing-instance-name instance-type name routing-options forwarding-table],  
[edit routing-options forwarding-table]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for routing instances added in Junos OS Release 8.3.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Statement introduced in Junos OS Release 11.3 for QFX Series switches.

NOTE: This feature is not supported on the EX4300 switch, even though it is available on the device.

Description

Control the operation of unicast reverse-path-forwarding check. This statement enables the RPF check to be used when routing is asymmetrical.

Options

active-paths—Consider only active paths during the unicast reverse-path check.

feasible-paths—Consider all feasible paths during the unicast reverse-path check.

Default: If you omit the **unicast-reverse-path** statement, only the active paths to a particular destination are considered.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Unicast RPF (On a Router)

wildcard-source (PIM RPF Selection)

Syntax

```
wildcard-source {  
  next-hop next-hop-address;  
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name protocols pim rpf-selection group group-address],  
[edit routing-instances routing-instance-name protocols pim rpf-selection prefix-list prefix-list-addresses]
```

Release Information

Statement introduced in Junos OS Release 10.4.

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

Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

view-level—To view this statement in the configuration.

control-level—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring PIM RPF Selection

Configuration Statements: IP Directed and Targeted Broadcast

IN THIS CHAPTER

- [targeted-broadcast](#) | 706
- [policy-statement](#) | 708

targeted-broadcast

Syntax (EX Series, MX Series, ACX Series)

```
targeted-broadcast {
  forward-and-send-to-re;
  forward-only;
}
```

Syntax (QFX Series, OCX1100, EX4600, NFX Series)

```
targeted-broadcast;
```

Hierarchy Level (EX Series, MX Series, ACX Series)

```
[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]
```

Hierarchy Level (QFX Series, OCX1100, EX4600, NFX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet],
[edit interfaces interface-range interface-range-name unit logical-unit-number family inet]
```

Release Information

Statement introduced in Junos OS Release 9.4 for EX Series switches.

Statement introduced in Junos OS Release 10.2.

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the IP packets destined for a Layer 3 broadcast address to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Default

When this statement is not included, broadcast packets are sent to the Routing Engine only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Targeted Broadcast | 229](#)

[Understanding Targeted Broadcast | 226](#)

policy-statement

Syntax

```

policy-statement policy-name {
  term term-name {
    from {
      as-path-unique-count count (equal | orhigher | orlower);
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      protocol protocol-name;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
      tag value;
      traffic-engineering;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
  then {
    aggregate-bandwidth;
    dynamic-tunnel-attributes dynamic-tunnel-attributes;
    limit-bandwidth limit-bandwidth;
    multipath-resolve;
    no-entropy-label-capability;
    prefix-segment {
      index index;
      node-segment;
    }
    priority (high | medium | low);
    resolution-map map-name;
  }
}

```

Hierarchy Level

```

[edit dynamic-profiles profile-name policy-options],
[edit logical-systems logical-system-name policy-options],

```

[edit policy-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

inet-mdt option introduced in Junos OS Release 10.0R2.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

route-target option introduced in Junos OS Release 12.2.

Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.

protocol and **traffic-engineering** options introduced in Junos OS Release 14.2.

no-entropy-label-capability option introduced in Junos OS Release 15.1.

priority and **tag value** options introduced in Junos OS Release 17.1.

as-path-unique-count option introduced in Junos OS Release 17.2R1.

prefix-segment option introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches, and QFX10000 switches.

multipath-resolve and **dynamic-tunnel-attributes** options introduced in Junos OS Release 17.3R1.

aggregate-bandwidth and **limit-bandwidth** ~~**limit-bandwidth**~~ options introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

l-isis and *l-ospf* keywords at the **protocol** option is introduced in Junos OS Release 19.1R1.

resolution-map statement introduced in Junos OS Release 19.2R1-S1 on MX and PTX Series routers.

lsp and **lsp-regex** options introduced in Junos OS Release 19.4R1.

Description

Define a routing policy, including subroutine policies.

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement** *policy-name* in alphabetical order, enter the **show policy-options** configuration command.

The statements are explained separately.

Options

actions—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

family *family-name*—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**. For traffic engineering, specify **traffic-engineering**.

NOTE: When **family** is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

from—(Optional) Match a route based on its source address.

as-path-unique-count *count* (*equal* | *orhigher* | *orlower*)—(Optional) Specify a number from 0 through 1024 to filter routes based on the number of unique autonomous systems (ASs) in the AS path. Specify the match condition for the unique AS path count.

aggregate-bandwidth—(Optional) Enable BGP to advertise aggregate outbound link bandwidth for load balancing.

dynamic-tunnel-attributes *dynamic-tunnel-attributes*—(Optional) Choose a set of defined dynamic tunnel attributes for forwarding traffic over V4oV6 tunnels.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

multipath-resolve *multipath-resolve*—(Optional) Enable the use of all paths for resolution over the specified prefix.

limit-bandwidth *limit-bandwidth*—(Optional) Specify the limit for advertised aggregate outbound link bandwidth for load balancing.

Range: 0 through 4,294,967,295 bytes

no-entropy-label-capability—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

priority (*high* | *medium* | *low*)—(Optional) Configure the priority for an IS-IS route to change the default order in which the routes are installed in the routing table, in the event of a network topology change.

policy subroutine-policy-name—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form `__.*-internal__`, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

prefix-list prefix-list-name—Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter prefix-list-name—Name of a prefix list to evaluate using qualifiers; **match-type** is the type of match, and **actions** is the action to take if the prefixes match.

protocol protocol-name—Name of the protocol used to control traffic engineering database import at the originating point.

Starting in Junos OS Release 19.1R1, you can specify options to match label IS-IS and label OSPF routes using the **l-isis** and **l-ospf** options, respectively. The **isis** options matches all IS-IS routes, excluding labelled IS-IS routes. The **ospf** option matches all OSPF routes, including OSPFv2, OSPFv3 and labelled OSPF routes.

resolution-map—(Optional) Set resolution map modes. A given resolution-map can be shared across multiple policy-statements.

route-filter destination-prefix match-type <actions>—(Optional) List of routes on which to perform an immediate match; **destination-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **destination-prefix** matches.

source-address-filter source-prefix match-type <actions>—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. **source-prefix** is the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **source-prefix** matches.

tag value—(Optional) A numeric value that identifies a route. You can tag certain routes to prioritize them over other routes. In the event of a network topology change, Junos OS updates these routes in the routing table before updating other routes with lower priority. You can also tag some routes to identify and reject them based on your requirement.

term term-name—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

dynamic-db

Understanding Source Packet Routing in Networking (SPRING)

Configuration Statements: ARP

IN THIS CHAPTER

- arp (Interfaces) | 715
- gratuitous-arp-reply | 719
- no-gratuitous-arp-request | 720
- proxy-arp | 721

arp (Interfaces)

Syntax

```
arp ip-address (mac | multicast-mac) mac-address publish;
```

```
arp {
  aging-timer minutes;
  gratuitous-arp-delay seconds;
  gratuitous-arp-on-ifup;
  interfaces {
    interface-name {
      aging-timer minutes;
    }
  }
  passive-learning;
  purging;
}
```

Syntax (EX Series)

```
arp {
  aging-timer minutes;
}
```

Hierarchy Level

```
[edit system]
```

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
```

```
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address]
```

NOTE: The **edit logical-systems** hierarchy is not available on QFabric systems.

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.

Description

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries mapping IP addresses to MAC addresses. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces interface-name unit logical-unit-number family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces interface-name unit logical-unit-number family inet]** hierarchy level.

NOTE: For EX-Series switches, set only the time interval between ARP updates.

Options

ip-address—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing **address** statement.

mac *mac-address*—MAC address to map to the IP 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**.

multicast-mac *mac-address*—Multicast MAC address to map to the IP address. Specify the multicast 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**.

publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.

NOTE: For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

aging-timer—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.

gratuitous-arp-delay— Configure a delay for gratuitous ARP requests at the system level. By default, Junos OS sends gratuitous ARP requests immediately after network-related configuration changes are made on an interface (for example, a VLAN ID, MAC address, IP address change, or Aggregated Ethernet deployment). This might lead to the Packet Forwarding Engine dropping some initial request packets if the configuration updates have not been fully processed. To avoid such request packets being dropped, you can configure a delay in gratuitous ARP requests.

Values:

- *seconds*—Configure the ARP request delay in seconds. We recommend configuring a value in the range of 3 through 6 seconds.

gratuitous-arp-on-ifup— Add this statement to the **[edit system arp]** hierarchy to configure Junos OS to automatically issue a gratuitous ARP announcement when an interface is online.

interfaces— Specify the ARP aging timer in minutes for a logical interface of family type **inet**.

Values: aging-timer *minutes*—Time between ARP updates, in minutes.

Default: 20

Range: 1 through 6,00,000

passive-learning— Configure backup VRRP routers or switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests. By default, the backup VRRP router drops these requests; therefore, if the master router fails, the backup router must learn all entries present in the ARP cache of the master router. Configuring passive learning reduces transition delay when the backup router is activated. Learning of ARP mappings (IP-to-MAC address) by backup VRRP routers or switches for hosts sending the requests is disabled unless this statement is configured.

purging— Purge obsolete ARP entries from the cache when an interface or link goes offline.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses | 238](#)

Configure ARP Learning and Aging Options

Junos OS Network Interfaces Library for Routing Devices

[Junos OS System Basics Configuration Guide](#) .

Adjusting the ARP Aging Timer

gratuitous-arp-reply

Syntax

```
(gratuitous-arp-reply | no-gratuitous-arp-reply);
```

Hierarchy Level

```
[edit interfaces interface-name],  
[edit interfaces interface-range interface-range-name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Enable processing of ARP updates received via gratuitous ARP reply messages.

Default

Updating of the ARP cache is disabled on all Ethernet interfaces.

Options

gratuitous-arp-reply—Update the ARP cache.

no-gratuitous-arp-reply—Do not update the ARP cache.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

no-gratuitous-arp-request

Syntax

```
no-gratuitous-arp-request;
```

Hierarchy Level

```
[edit interfaces interface-name],  
[edit interfaces interface-range interface-name]
```

Release Information

Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs).

Default

Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring IRB Interfaces on Switches*

proxy-arp

Syntax

```
proxy-arp (restricted | unrestricted);
```

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 9.6 for EX Series switches.

restricted added in Junos OS Release 10.0 for EX Series switches.

Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description

For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.

NOTE: You must configure the IP address and the **inet** family for the interface when you enable proxy ARP.

Default

Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.

Options

- **none**—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
- **restricted**—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.
- **unrestricted**—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.

Default: **unrestricted**

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Restricted and Unrestricted Proxy ARP | 242](#)

Configuring Proxy ARP on Switches

Example: Configuring Proxy ARP on an EX Series Switch

[Configuring Gratuitous ARP | 243](#)

Configuration Statements: Resilient Hashing

IN THIS CHAPTER

- [ecmp-resilient-hash](#) | 724
- [enhanced-hash-key](#) | 725
- [hash-key \(Forwarding Options\)](#) | 731
- [hash-mode](#) | 733
- [hash-seed](#) | 735
- [inet \(enhanced-hash-key\)](#) | 736
- [inet6 \(enhanced-hash-key\)](#) | 739
- [ipv6-flow-label](#) | 742
- [resilient-hash](#) | 743

ecmp-resilient-hash

Syntax

```
ecmp-resilient-hash;
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description

Enable resilient hashing for ECMP groups, to minimize remapping of destination paths.



NOTE:

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Resilient Hashing for LAGs/ECMP Groups](#) | 248

enhanced-hash-key

List of Syntax

[Syntax \(EX Series\) on page 725](#)

[Syntax \(QFX5000 Line of Switches\) on page 726](#)

[Syntax \(QFX10000 Series Switches\) on page 728](#)

Syntax (EX Series)

```
enhanced-hash-key {
  ecmp-resilient-hash;
  fabric-load-balance {
    flowlet {
      inactivity-interval interval;
    }
    per-packet;
  }
  hash-mode {
    layer2-header;
    layer2-payload;
  }
  family inet {
    no-ipv4-destination-address;
    no-ipv4-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-protocol;
    vlan-id;
  }
  family inet6 {
    no-ipv6-destination-address;
    no-ipv6-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-next-header;
    vlan-id;
  }
  layer2 {
    no-destination-mac-address;
    no-ether-type;
    no-source-mac-address;
    vlan-id;
  }
}
```

Syntax (QFX5000 Line of Switches)

```

enhanced-hash-key {
  ecmp-dlb {
    assigned-flow;
    per-packet;
    flowlet inactivity-interval;
    ether-type (ipv4|ipv6|mpls);
  }
  ecmp-resilient-hash;
  fabric-load-balance {
    flowlet {
      inactivity-interval interval;
    }
    per-packet;
  }
  hash-mode {
    layer2-header;
    layer2-payload;
    gtp-header-offset offset-value;
  }
  hash-parameters {
    ecmp {
      function {
        (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
      }
      offset offset;
      preprocess;
    }
    lag {
      function {
        (crc16-bisync | crc16-ccitt | crc32-hi | crc32-lo);
      }
      offset offset;
      preprocess;
    }
  }
}
family inet {
  gtp-tunnel-endpoint-identifier;
  no-ipv4-destination-address;
  no-ipv4-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-protocol;
  vlan-id;
}

```

```
}  
family inet6 {  
    no-ipv6-destination-address;  
    no-ipv6-source-address;  
    no-l4-destination-port;  
    no-l4-source-port;  
    no-next-header;  
    vlan-id;  
}  
layer2 {  
    no-destination-mac-address;  
    no-ether-type;  
    no-source-mac-address;  
    vlan-id;  
}  
}
```

Syntax (QFX10000 Series Switches)

```

enhanced-hash-key {
  hash-seed seed-value;
  family inet {
    gtp-tunnel-endpoint-identifier;
    no-ipv4-destination-address;
    no-ipv4-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-incoming-port;
  }
  family inet6 {
    gtp-tunnel-endpoint-identifier;
    ipv6-flow-label;
    no-ipv6-destination-address;
    no-ipv6-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-incoming-port;
  }
  layer2 {
    destination-mac-address
    inner-vlan-id;
    no-ether-type;
    no-vlan-id;
    source-mac-address;
  }
  no-mpls;
  gre {
    key;
    protocol;
  }
  vxlan-vnid
}

```

Hierarchy Level

[edit forwarding-options]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

The **fabric-load-balance** statement introduced in Junos OS Release 14.1X53-D10.

The **fabric-load-balance** statement deprecated starting in Junos OS Releases 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1.

The **hash-seed** statement introduced in Junos OS Release 15.1X53-D30.

The **ecmp-dlb** statement introduced in Junos OS Release 19.4R1 for QFX5120-32C and QFX5120-48Y switches.

Description

Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

NOTE: Starting in Junos OS Release 14.1X53-D46, 15.1R7, 16.1R6, 17.1R3, 17.2R2, 17.3R2, and 17.4R1, the ALB feature is deprecated. If **fabric-load-balance** is enabled in the configuration for a VCF, delete the configuration item upon upgrading Junos OS.

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

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

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

On QFX10000 Series switches, you can configure the hash seed for load balancing. By default, the QFX10000 Series switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the **hash-seed** statement at the [edit forwarding-options enhanced-hash-key] hierarchy level. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system generates a hash seed value based on the system MAC address.

The remaining statements are explained separately. See [CLI Explorer](#).

Starting in Junos OS Release 18.4R1, symmetric hashing is supported on the QFX10000 Series switches. You configure the **no-incoming-port** option under the [edit forwarding-options [enhanced-hash-key](#)] hierarchy. By default, Dynamic IP (DIP), SIP, Layer 4 source and destination ports, and the incoming port are used for hashing. You can only configure symmetric hashing at the global level.

Starting in Junos OS Release 19.4R1, the dynamic load balancing on ECMP is supported on QFX5120-32C and QFX5120-48Y switches. You can configure the **ecmp-dlb** option under the [edit forwarding-options [enhanced-hash-key](#)] hierarchy. Refer "[Dynamic Load Balancing](#)" on [page 372](#) for more details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) 349
Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic 342
Understanding Passive Monitoring
Understanding Per-Packet Load Balancing 257
show forwarding-options enhanced-hash-key 818

hash-key (Forwarding Options)

Syntax

```
hash-key {
  family {
    inet {
      layer-3;
      layer-4;
      inner-vlan-id;
      outer-vlan-id;
    }
  }
  multiservice {
    source-mac;
    destination-mac;
    payload {
      ip {
        layer3-only;
        layer-3 (source-ip-only | destination-ip-only);
        layer-4;
        inner-vlan-id;
        outer-vlan-id;
      }
    }
  }
}
```

Hierarchy Level

```
[edit forwarding-options]
[edit chassis fpc slot-number pic pic-number]
```

Release Information

Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

Description

(QFX10000 switches only) Select which packet header data to use for per-flow load balancing.

NOTE: You can configure either Layer 3 or Layer 4 load balancing, or both at the same time.

NOTE: On I chip platforms, an unknown Layer 4 header is excluded from load-balance hashing to avoid undesired packet reordering.

Options

- **inet**—IPv4 address family.
- **layer-3**—Incorporate Layer 3 data into the hash key.
- **layer-4**—Incorporate Layer 4 data into the hash key.
- **outer-vlan-id**—Include outer VLAN ID information in the hash key.
- **inner-vlan-id**—Include inner VLAN ID information in the hash key.
- **payload**—Incorporate payload data into the hash key.
- **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key.
- **layer-3-only**—Include only Layer 3 IP information.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

hash-mode

Syntax

```
hash-mode {
    layer2-header;
    layer2-payload;
    gtp-header-offset offset-value;
}
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key]
```

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

Statement is not supported on QFX10002 and QFX 10008 switches.

Hash mode option *gtp-header-offset* introduced in Junos OS Release 17.3R3 for QFX5000 line of switches.

Description

Select the mode for the hashing algorithm.

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

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

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

The hash mode that is set using this statement determines which fields are inspected by the hashing algorithm. You must set the hash mode to **layer2-payload** if you want the hashing algorithm to inspect fields in the Layer 2 payload when making hashing decisions. You must set the hash mode to **layer2-header** if you want the hashing algorithm to inspect fields in the Layer 2 header when making hashing decisions.

If the hash mode is set to **layer2-payload**, you can set the fields used by the hashing algorithm to hash IPv4 traffic using the **set forwarding-options enhanced-hash-key inet** statement. You can set the fields used by the hashing algorithm to hash IPv6 traffic using the **set forwarding-options enhanced-hash-key inet6** statement.

If the hash mode is set to **layer2-header**, you can set the fields that the hashing algorithm inspects in the Layer 2 header using the **set forwarding-options enhanced-hash-key layer2** statement.

Default

layer2-payload

Options

layer-2-payload—Set the hashing algorithm to use fields in the Layer 2 payload to make hashing decisions.

layer-2-header—Set the hashing algorithm to use fields in the Layer 2 header to make hashing decisions.

gtp-header-offset—Set proper offset value for the header based on the different patterns.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) 349
Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic 342
enhanced-hash-key 725
inet 443
inet6 446
layer2 670
gtp-header-offset 745

hash-seed

Syntax

```
hash-seed seed-value;
```

Hierarchy Level

```
[edit forwarding-options enhanced-hash-key ]
```

Release Information

Statement introduced in Junos OS Release 15.1X53-D30 on QFX Series devices.

Description

Configure a hash seed for load-balancing functions.

By default, the QFX10002 and QFX10008 switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the **hash-seed** statement at the **[edit forwarding-options enhanced-hash-key]** hierarchy. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system will generate a hash seed value based on the system MAC address.

NOTE: The **fabric-load-balance** and **user-defined-fields** statements are not supported at the **[edit forwarding-options enhanced-hash-key]** hierarchy level.

Options

hash-seed*seed-value*—A hash seed value, in the range from 0 to 4294967295.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

[Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\)](#) | 349

[show forwarding-options enhanced-hash-key](#) | 818

inet (enhanced-hash-key)

List of Syntax

[Syntax \(EX Series and QFX5100 Switch\) on page 736](#)

[Syntax \(QFX10000 Series Switches\) on page 736](#)

Syntax (EX Series and QFX5100 Switch)

```
inet {
  gtp-tunnel-endpoint-identifier;
  no-ipv4-destination-address;
  no-ipv4-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-protocol;
  vlan-id;
}
```

Syntax (QFX10000 Series Switches)

```
inet {
  gtp-tunnel-endpoint-identifier;
  no-ipv4-destination-address;
  no-ipv4-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-incoming-port;
}
```

Hierarchy Level

[edit forwarding-options [enhanced-hash-key](#) family]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

Statement introduced in Junos OS Release 15.1X53-D30 on QFX10000 Series Switches.

Description

Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.

When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the **inet** statement and uses the information in the fields to decide how to place

traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to **layer2-payload**. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the **set forwarding-options enhanced-hash-key hash-mode layer2-payload** statement.

Default

The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:

- IP destination address
- IP source address
- Layer 4 destination port
- Layer 4 source port
- Protocol

Options

no-ipv4-destination-address—Exclude the IPv4 destination address field from the hashing algorithm.

no-ipv4-source-address—Exclude the IPv4 source address field from the hashing algorithm.

no-l4-destination-port—Exclude the Layer 4 destination port field from the hashing algorithm.

no-l4-source-port—Exclude the Layer 4 source port field from the hashing algorithm.

no-protocol—Exclude the protocol field from the hashing algorithm.

no-incoming-port—Exclude the incoming port number from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

NOTE: The **vlan-id** option is not supported and should not be configured on a Virtual Chassis or Virtual Chassis Fabric (VCF) that contains any of the following switches as members: EX4300, EX4600, QFX3500, QFX3600, QFX5100, or QFX5110 switches.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)	349
Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic	342
<i>Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)</i>	
Understanding Per-Packet Load Balancing	257
hash-seed	735
enhanced-hash-key	725
hash-mode	733
inet6	446

inet6 (enhanced-hash-key)

List of Syntax

[Syntax \(EX Series and QFX5100 Switch\) on page 739](#)

[Syntax \(QFX10000 Series Switches\) on page 739](#)

Syntax (EX Series and QFX5100 Switch)

```
inet6 {
  no-ipv6-destination-address;
  no-ipv6-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-next-header;
  vlan-id;
}
```

Syntax (QFX10000 Series Switches)

```
inet6 {
  gtp-tunnel-endpoint-identifier;
  ipv6-flow-label;
  no-ipv6-destination-address;
  no-ipv6-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-incoming-port;
}
```

Hierarchy Level

[edit forwarding-options [enhanced-hash-key](#) family]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 on EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 on QFX Series devices.

Statement introduced in Junos OS Release 15.1X53-D30 on QFX10000 Series switches.

Description

Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.

When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place

traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the **set forwarding-options enhanced-hash-key hash-mode layer2-payload** statement.

Default

The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:

- IP destination address
- IP source address
- Layer 4 destination port
- Layer 4 source port
- Next header

Options

no-ipv6-destination-address—Exclude the IPv6 destination address field from the hashing algorithm.

no-ipv6-source-address—Exclude the IPv6 source address field from the hashing algorithm.

no-l4-destination-port—Exclude the Layer 4 destination port field from the hashing algorithm.

no-l4-source-port—Exclude the Layer 4 source port field from the hashing algorithm.

no-incoming-port—Exclude the incoming port number from the hashing algorithm.

no-next-header—Exclude the Next Header field from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) 349
Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic 342
<i>Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)</i>
Understanding Per-Packet Load Balancing 257
hash-seed 735
enhanced-hash-key 725
hash-mode 733
inet 443

ipv6-flow-label

Syntax

```
ipv6-flow-label;
```

Hierarchy Level

For QFX Series:

```
[edit forwarding-options enhanced-hash-key inet6]
```

Hierarchy Level

For PTX Series:

```
[edit forwarding-options hash-key family inet6 layer-3]
```

Release Information

Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and 10008 switches..

Statement introduced in Junos OS Release 18.3R1 for PTX series routers.

Description

Enable IPv6 packet flow labels for hash calculations and load balancing based on the flow label of the IPv6 header.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)

[show forwarding-options enhanced-hash-key](#) | 818

resilient-hash

Syntax

```
resilient-hash;
```

Hierarchy Level

```
[edit interfaces aex aggregated-ether-options]]
```

Release Information

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description

Enable resilient hashing for a LAG to minimize remapping of destination paths.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Resilient Hashing for LAGs/ECMP Groups](#) | 248

Configuration Statements: Generic Routing Encapsulation (GRE)

IN THIS CHAPTER

- gtp-header-offset | 745
- gtp-tunnel-endpoint-identifier | 747
- source | 749
- ttl | 750
- tunnel | 751
- allow-fragmentation | 752
- copy-tos-to-outer-ip-header | 753
- do-not-fragment | 754
- destination (Tunnels) | 755
- family | 756
- routing-instance | 758
- source | 759
- tunnel | 760
- tunnel-port | 761
- unit (Interfaces) | 762

gtp-header-offset

Syntax

```
gtp-header-offset offset-value
```

Hierarchy Level (QFX5000 line of switches)

```
[edit forwarding-options enhanced-hash-key hash-mode ]
```

Release Information

Statement introduced in Junos OS Release 17.3R3 for QFX5000 line of switches.

Description

You must only provide a decimal value for *gtp-header-offset* command. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary like, 0x30, 0x28, and so on. In this case, you need to convert it to a proper offset value as follows:

For Example:

- If the offset value is 0x32, the equivalent decimal value is calculated as $(16*3+2)$ which is equal to 50.
- If the offset value is 0x30, the equivalent decimal value is calculated as $(16*3+0)$ which is equal to 48.

After setting *gtp-header-offset* to a proper offset value, run *gtp-tunnel-endpoint-identifier* command to enable GTP hashing. Refer [gtp-tunnel-endpoint-identifier](#) for more details. For example:

- If the IPv6 traffic has 0x32 as offset value, then use **set forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier** command directly.
- If the IPv6 traffic has 0x30 as offset value, then you must set proper offset value after enabling GTP hashing as follows:

```
set forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier
```

```
set forwarding-options enhanced-hash-key hash-mode gtp-header-offset 48
```

- If the IPv4 traffic has 0x32 as offset value, then use **set forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier** command directly.
- If the IPv4 traffic has 0x30 as offset value, then you must set proper offset value after enabling GTP hashing as follows:

```
set forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier
```

```
set forwarding-options enhanced-hash-key hash-mode gtp-header-offset 48
```

NOTE: Both IPv4 and IPv6 share same GTP header offset value. If there are two different traffic streams, that is, if IPv4 has value X as offset value and IPv6 has Y as offset value, then hashing will be successful only for one of the traffic stream.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [gtp-tunnel-endpoint-identifier](#) | 747

gtp-tunnel-endpoint-identifier

Syntax

```
gtp-tunnel-endpoint-identifier;
```

Hierarchy Level

```
[edit forwarding-options hash-key family inet layer-4],
```

```
[edit forwarding-options hash-key family inet6 layer-4]
```

Hierarchy Level (QFX5000 line of switches)

```
[edit forwarding-options enhanced-hash-key family inet]
```

Hierarchy Level (QFX10000 line of switches)

```
[edit forwarding-options enhanced-hash-key family inet],
```

```
[edit forwarding-options enhanced-hash-key family inet6]
```

Release Information

Statement introduced in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers with third-generation FPCs.

Statement introduced in Junos OS Release 15.1F6 and 16.1R2 for PTX3000 routers with third-generation FPCs.

Statement introduced in Junos OS Release 17.3R3 for the QFX5000 line of switches.

Statement introduced in Junos OS Release 19.1R1 for the QFX10000 line of switches.

Description

When you configure **gtp-tunnel-endpoint-identifier**, the hash calculation of IPv4 or IPv6 packets are included in the GPRS tunneling protocol–tunnel endpoint ID (GTP-TEID) field hash calculations.

NOTE: The **gtp-tunnel-endpoint-identifier** configuration statement is supported on PTX Series routers only when network services is set to **enhanced-mode**. For more information, see **enhanced-mode**.

On the QFX5000 and QFX10000 lines of switches, if the **gtp-tunnel-endpoint-identifier** statement is configured, the default Layer 4 port 2152 (and 2123 as well on QFX5000) is set to use along with the default first byte 0x32.

(On the QFX5000 line of switches only) In most cases, configuring the **gtp-tunnel-endpoint-identifier** statement on QFX5000 switches is sufficient for enabling GTP hashing. After you have enabled GTP hashing, if GTP hashing does not work, we recommend that you capture the packets using relevant tools and identify the offset value. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary to say 0x30, 0x28, and so on. In these cases, use the **gtp-header-offset** statement to set a proper offset value. Once the header offset value is resolved, run the **gtp-tunnel-endpoint-identifier** command for enabling GTP hashing successfully. Refer to [gtp-header-offset](#) for more details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>hash-key</i>
Understanding Per-Packet Load Balancing 257
<i>Configuring Per-Packet Load Balancing</i>
gtp-header-offset 745

source

Syntax

```
source source-address;
```

Hierarchy Level (EX, NFX, OCX1100 and QFX Series)

```
[edit interfaces interface-name unit logical-unit-number tunnel]
```

Hierarchy Level (M-series and T-series)

```
[edit interfaces interface-name unit logical-unit-number tunnel address],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number tunnel address]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the source address of the tunnel.

Default

If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.

Options

source-address—Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Tunnel Services Overview

multicast-only

[primary \(Address on Interface\) | 494](#)

Junos OS Services Interfaces Library for Routing Devices

ttl

Syntax

```
ttl value;
```

Hierarchy Level

```
[edit interfaces interface-name unit number tunnel]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 17.1 for ACX Series routers.

Description

Set the time-to-live value bit in the header of the outer IP packet.

Options

value—Time-to-live value.

Range: 0 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

Tunnel Services Overview

tunnel

Syntax

```
tunnel {  
  destination destination-address;  
  source source-address;  
  ttl number;  
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic.

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 Generic Routing Encapsulation Tunneling](#) | 255

allow-fragmentation

Syntax

```
allow-fragmentation;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel],  
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number tunnel]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

Statement introduced in Junos OS Release 19.3 for MPC10E line card.

Description

For a generic routing encapsulation (GRE) tunnel, enable fragmentation of GRE-encapsulated packets whose size exceeds the maximum transmission unit (MTU) value of a link that the packet passes through. The don't fragment (DF) bit is not set in the outer IP header of GRE-encapsulated packets.

To enable the reassembly of fragmented GRE-encapsulated packets on GRE tunnel interfaces at the endpoint of the GRE tunnel, include the *reassemble-packets* statement for the interface.

NOTE: The *reassemble-packets* statement is not supported on MPC10E line card in Junos OS Release 19.3.

Default

If you do not include the **allow-fragmentation** statement, fragmentation of GRE-encapsulated packets is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

reassemble-packets

147526Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation

copy-tos-to-outer-ip-header

Syntax

```
copy-tos-to-outer-ip-header;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number],  
[edit interfaces gre unit logical-unit-number],  
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number],  
[edit logical-systems logical-system-name interfaces gre unit logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 8.2.

Support for GRE interfaces for Generalized MPLS (GMPLS) introduced in Junos OS Release 12.3R7.

Description

For GRE tunnel interfaces and GRE interfaces for GMPLS control channels only, enable the inner IP header's ToS bits to be copied to the outer IP packet header for traffic originating in the Routing Engine.

To verify that this option is enabled at the interface level, use the **show interfaces *interface-name* detail** command.

Default

If you omit this statement, the ToS bits in the outer IP header are set to 0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

copy-tos-to-outer-ip-header-transit

force-control-packets-on-transit-path

do-not-fragment

Syntax

```
do-not-fragment;
```

Hierarchy Level

```
[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel],  
[edit logical-systems logical-system-name interfaces gr-fpc/pic/port unit logical-unit-number tunnel]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

Statement introduced in Junos OS Release 19.3 for MPC10E line card.

Description

For a generic routing encapsulation (GRE) tunnel, disable fragmentation of GRE-encapsulated packets. This sets the do-not-fragment (DF) bit in the outer IP header of the GRE-encapsulated packets so that they do not get fragmented anywhere in the path. When the size of a GRE-encapsulated packet is greater than the MTU of a link that the packet passes through, the GRE-encapsulated packet is dropped.

Default

By default, fragmentation of GRE-encapsulated packets is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[allow-fragmentation](#) | 752

reassemble-packets

Enabling Fragmentation and Reassembly on Packets After GRE-Encapsulation

Junos OS Services Interfaces Library for Routing Devices

destination (Tunnels)

Syntax

```
destination address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet address address],
[edit interfaces interface-name unit logical-unit-number family inet unnumbered-address interface-name],
[edit interfaces interface-name unit logical-unit-number tunnel],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address
  address],
[edit logical-systems logical-system-name 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 tunnel]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.

Options

address—Address of the remote side of the connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Interface Address | 76](#)

[Configuring Generic Routing Encapsulation Tunneling | 255](#)

Junos OS Services Interfaces Library for Routing Devices

family

Syntax

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

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

Description

Configure protocol family information for the logical interface.

Options

family—Protocol family:

- **ccc**—Circuit cross-connect protocol suite
- **inet**—IP version 4 (IPv4)
- **inet6**—IP version 6 (IPv6)
- **iso**—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite
- **mlfr-end-to-end**—Multilink Frame Relay FRF.15
- **mlfr-uni-nni**—Multilink Frame Relay FRF.16
- **multilink-ppp**—Multilink Point-to-Point Protocol
- **mpls**—MPLS
- **tcc**—Translational cross-connect protocol suite
- **tnp**—Trivial Network Protocol
- **vpis**—Virtual private LAN service

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

Link and Multilink Services Interfaces User Guide for Routing Devices

Junos OS Network Interfaces Library for Routing Devices

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.

Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

Description

Specify the destination routing instance that points to the routing table containing the tunnel destination address.

Default

The default Internet routing table **inet.0**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring Tunnel Interfaces for Routing Table Lookup*

source

Syntax

```
source source-address;
```

Hierarchy Level (EX, NFX, OCX1100 and QFX Series)

```
[edit interfaces interface-name unit logical-unit-number tunnel]
```

Hierarchy Level (M-series and T-series)

```
[edit interfaces interface-name unit logical-unit-number tunnel address],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number tunnel address]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the source address of the tunnel.

Default

If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.

Options

source-address—Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Tunnel Services Overview

multicast-only

tunnel

Syntax

```
tunnel {  
  destination destination-address;  
  source source-address;  
  ttl number;  
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Release Information

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic.

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

tunnel-port

Syntax

```
tunnel-port port-number tunnel-services;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the port number for generic routing encapsulation (GRE) tunneling.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Generic Routing Encapsulation Tunneling](#) | 255

unit (Interfaces)

Syntax

```

unit logical-unit-number {
    peer-unit unit-number;
    reassemble-packets;
    tunnel {
        allow-fragmentation;
        backup-destination address;
        destination destination-address;
        do-not-fragment;
        key number;
        routing-instance {
            destination routing-instance-name;
        }
        source source-address;
        ttl number;
    }
}

```

Hierarchy Level

```

[edit interfaces interface-name],
[edit logical-systems logical-system-name interfaces interface-name]

```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.

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 16,384

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

Junos OS Network Interfaces Library for Routing Devices for other statements that do not affect services interfaces.

Configuration Statements: Flexible Ethernet Services Encapsulation

IN THIS CHAPTER

- [encapsulation | 765](#)
- [encapsulation \(Logical Interface\) | 772](#)
- [flexible-vlan-tagging | 777](#)

encapsulation

List of Syntax

[Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series on page 765](#)

[Syntax for Physical Interfaces: SRX Series on page 765](#)

[Syntax for Logical Interfaces: SRX Series on page 765](#)

Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

```
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc | ethernet-bridge | ethernet-ccc
| ethernet-over-atm | ethernet-tcc | ethernet-vpls | ethernet-vpls-fr | ether-vpls-over-atm-llc | ethernet-vpls-ppp
| extended-frame-relay-ccc | extended-frame-relay-ether-type-tcc | extended-frame-relay-tcc |
extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services |
flexible-frame-relay | frame-relay | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc |
frame-relay-port-ccc | frame-relay-tcc | generic-services | multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc
| vlan-ccc | vlan-vci-ccc | vlan-vpls);
```

Syntax for Physical Interfaces: SRX Series

```
encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc | ethernet-vpls |
extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc |
extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services | frame-relay-port-ccc | vlan-ccc | vlan-vpls);
```

Syntax for Logical Interfaces: SRX Series

```
encapsulation ( dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-bridge | vlan-ccc | vlan-tcc | vlan-vpls );
```

Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

```
[edit interfaces interface-name],
[edit interfaces rlsq number:number]
```

Logical Interfaces

```
[edit interfaces interface-name unit logical-unit-number ]
```

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 \(CLI Procedure\)](#)

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.

Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms. See *Configuring VLAN TCC Encapsulation* for more details. Non-ethernet media types, SONET and ATM interfaces are only supported. It is expected that the user will have the member links of aggregated ethernet with supported hardware for configuring VLAN TCC encapsulation and no commit check is performed externally for the aggregated ethernet (AE) interfaces.

Options

atm-ccc-cell-relay—Use ATM cell-relay encapsulation.

atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

atm-cisco-nlpid—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-ppp-llc—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.

atm-ppp-vc-mux—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.

atm-snap—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.

NOTE: The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.

ether-vpls-over-ppp—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE router over a time-division multiplexing (TDM) link. This encapsulation type enables the PE router to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

ppp-over-ether-over-atm-llc—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

vxlan—Use VXLAN data plane encapsulation for EVPN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release History Table

Release	Description
20.1R1	Starting in Junos OS Release 20.1R1, aggregated ethernet interfaces supports VLAN TCC (Translational cross-connect) encapsulation on MX series platforms.

RELATED DOCUMENTATION

Configuring Layer 2 Switching Cross-Connects Using CCC

Configuring the Encapsulation for Layer 2 Switching TCCs

Configuring Interface Encapsulation on Logical Interfaces

Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects

Circuit and Translational Cross-Connects Overview

Identifying the Access Concentrator

Configuring ATM Interface Encapsulation

Configuring VLAN and Extended VLAN Encapsulation

Configuring ATM-to-Ethernet Interworking

Configuring Interface Encapsulation on PTX Series Packet Transport Routers

Configuring CCC Encapsulation for Layer 2 VPNs

Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits

Configuring ATM for Subscriber Access

Understanding CoS on ATM IMA Pseudowire Interfaces Overview

Configuring Policing on an ATM IMA Pseudowire

flexible-vlan-tagging

Syntax

```
flexible-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces aex],  
[edit interfaces ge-fpc/pic/port],  
[edit interfaces et-fpc/pic/port],  
[edit interfaces ps0],  
[edit interfaces xe-fpc/pic/port]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Support for aggregated Ethernet added in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description

Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.

This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP.

This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series and QFX 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 VLAN Tagging

Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers

Configuring Double-Tagged VLANs on Layer 3 Logical Interfaces

Operational Commands

IN THIS CHAPTER

- [Common Output Fields Description | 779](#)
- [clear interfaces statistics | 788](#)
- [monitor interface | 790](#)
- [request diagnostics tdr | 806](#)
- [request chassis system-mode | 808](#)
- [Show chassis system-mode | 810](#)
- [show diagnostics tdr | 812](#)
- [show forwarding-options enhanced-hash-key | 818](#)
- [show interfaces \(Discard\) | 825](#)
- [show interfaces | 832](#)
- [show interfaces \(Serial\) | 941](#)
- [show interfaces diagnostics optics | 959](#)
- [show interfaces extensive | 990](#)
- [show interfaces fabric | 1047](#)
- [show interfaces ge | 1075](#)
- [show interfaces \(GRE\) | 1092](#)
- [show interfaces irb | 1104](#)
- [show interfaces mc-ae | 1113](#)
- [show interfaces me0 | 1117](#)
- [show interfaces queue | 1127](#)
- [show interfaces queue fabric | 1183](#)
- [show interfaces xe | 1213](#)
- [show interfaces xle | 1237](#)
- [show interfaces statistics fabric | 1260](#)
- [show interfaces vlan | 1285](#)
- [show lacp interfaces | 1301](#)
- [show lacp statistics interfaces \(View\) | 1308](#)
- [show redundant-trunk-group | 1310](#)

- [show uplink-failure-detection | 1312](#)
- [show virtual-chassis vc-port diagnostics optics | 1315](#)
- [test interface restart-auto-negotiation | 1334](#)

Common Output Fields Description

This chapter explains the content of the output fields, which appear in the output of most **show interfaces** commands.

Damping Field

For the physical interface, the Damping field shows the setting of the following damping parameters:

- **half-life**—Decay half-life. The number of seconds after which the accumulated interface penalty counter is reduced by half if the interface remains stable.
- **max-suppress**—Maximum hold-down time. The maximum number of seconds that an interface can be suppressed irrespective of how unstable the interface has been.
- **reuse**—Reuse threshold. When the accumulated interface penalty counter falls below this number, the interface is no longer suppressed.
- **suppress**—Cutoff (suppression) threshold. When the accumulated interface penalty counter exceeds this number, the interface is suppressed.
- **state**—Interface damping state. If damping is enabled on an interface, it is suppressed during interface flaps that match the configured damping parameters.

Destination Class Field

For the logical interface, the **Destination class** field provides the names of destination class usage (DCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980

	(889)	(597762)
bronze		0		0
	(0)	(0)
silver		0		0
	(0)	(0)

Enabled Field

For the physical interface, the **Enabled** field provides information about the state of the interface, displaying one or more of the following values:

- **Administratively down, Physical link is Down**—The interface is turned off, and the physical link is inoperable and cannot pass packets even when it is enabled. To change the interface state to **Enabled**, use the following command:

```
user@host# set interfaces interface enable
```

Manually verify the connections to bring the physical link up.

- **Administratively down, Physical link is Up**—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. To change the interface state to **Enabled**, use the following command:

```
user@host# set interfaces interface enable
```

- **Enabled, Physical link is Down**—The interface is turned on, but the physical link is inoperable and cannot pass packets. Manually verify the connections to bring the physical link up.
- **Enabled, Physical link is Up**—The interface is turned on, and the physical link is operational and can pass packets.

Filters Field

For the logical interface, the **Filters** field provides the name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is **Filters: Input: *filter-name*** and **Filters: Output: *filter-name***. For example:

```
Filters: Input: sample-all
Filters: Output: cp-ftp
```

Flags Fields

IN THIS SECTION

- [Addresses, Flags Field | 781](#)
- [Device Flags Field | 782](#)
- [Family Flags Field | 782](#)
- [Interface Flags Field | 783](#)
- [Link Flags Field | 784](#)
- [Logical Interface Flags Field | 784](#)

The following sections provide information about flags that are specific to interfaces:

Addresses, Flags Field

The **Addresses, Flags** field provides information about the addresses configured for the protocol family on the logical interface and displays one or more of the following values:

- **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status
- **Is-Default**—The default address of the router used as the source address by SNMP, ping, traceroute, and other network utilities.
- **Is-Preferred**—The default local address for packets originating from the local router and sent to destinations on the subnet.
- **Is-Primary**—The default local address for broadcast and multicast packets originated locally and sent out the interface.
- **Preferred**—This address is a candidate to become the preferred address.
- **Primary**—This address is a candidate to become the primary address.
- **Trunk**—Interface is a trunk.
- **Trunk, Inter-Switch-Link**—Interface is a trunk, and InterSwitch Link protocol (ISL) is configured on the trunk port of the primary VLAN in order to connect the routers composing the PVLAN to each other.

Device Flags Field

The **Device flags** field provides information about the physical device and displays one or more of the following values:

- **ASIC Error**—Device is down because of ASIC wedging and due to which PFE is disabled.
- **Down**—Device has been administratively disabled.
- **Hear-Own-Xmit**—Device receives its own transmissions.
- **Link-Layer-Down**—The link-layer protocol has failed to connect with the remote endpoint.
- **Loopback**—Device is in physical loopback.
- **Loop-Detected**—The link layer has received frames that it sent, thereby detecting a physical loopback.
- **No-Carrier**—On media that support carrier recognition, no carrier is currently detected.
- **No-Multicast**—Device does not support multicast traffic.
- **Present**—Device is physically present and recognized.
- **Promiscuous**—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.
- **Quench**—Transmission on the device is quenched because the output buffer is overflowing
- **Recv-All-Multicasts**—Device is in multicast promiscuous mode and therefore provides no multicast filtering.
- **Running**—Device is active and enabled.

Family Flags Field

The **Family flags** field provides information about the protocol family on the logical interface and displays one or more of the following values:

- **DCU**—Destination class usage is enabled.
- **Dest-route-down**—The software detected that the link is down and has stopped forwarding the link's interface routes.
- **Down**—Protocol is inactive.
- **Is-Primary**—Interface is the primary one for the protocol.
- **Mac-Validate-Loose**—Interface is enabled with loose MAC address validation.
- **Mac-Validate-Strict**—Interface is enabled with strict MAC address validation.
- **Maximum labels**—Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.
- **MTU-Protocol-Adjusted**—The effective MTU is not the configured value in the software.
- **No-Redirects**—Protocol redirects are disabled.

- **Primary**—Interface can be considered for selection as the primary family address.
- **Protocol-Down**—Protocol failed to negotiate correctly.
- **SCU-in**—Interface is configured for source class usage input.
- **SCU-out**—Interface is configured for source class usage output.
- **send-bcast-packet-to-re**—Interface is configured to forward IPv4 broadcast packets to the Routing Engine.
- **targeted-broadcast**—Interface is configured to forward IPv4 broadcast packets to the LAN interface and the Routing Engine.
- **Unnumbered**—Protocol family is configured for unnumbered Ethernet. An unnumbered Ethernet interface borrows an IPv4 address from another interface, which is referred to as the donor interface.
- **Up**—Protocol is configured and operational.
- **uRPF**—Unicast Reverse Path Forwarding is enabled.

Interface Flags Field

The **Interface flags** field provides information about the physical interface and displays one or more of the following values:

- **Admin-Test**—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Hardware-Down**—Interface is nonfunctional or incorrectly connected.
- **Link-Layer-Down**—Interface keepalives have indicated that the link is incomplete.
- **No-Multicast**—Interface does not support multicast traffic.
- **No-receive No-transmit**—Passive monitor mode is configured on the interface.
- **OAM-On-SVLAN**—(MX Series routers with MPC/MIC interfaces only) Interface is configured to propagate the Ethernet OAM state of a static, single-tagged service VLAN (S-VLAN) on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interface to a dynamic or static double-tagged customer VLAN (C-VLAN) that has the same S-VLAN (outer) tag as the S-VLAN.
- **Point-To-Point**—Interface is point-to-point.
- **Pop all MPLS labels from packets of depth**—MPLS labels are removed as packets arrive on an interface that has the **pop-all-labels** statement configured. The depth value can be one of the following:
 - **1**—Takes effect for incoming packets with one label only.
 - **2**—Takes effect for incoming packets with two labels only.
 - **[1 2]**—Takes effect for incoming packets with either one or two labels.

- **Promiscuous**—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
- **Recv-All-Multicasts**—Interface is in multicast promiscuous mode and provides no multicast filtering.
- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Link Flags Field

The **Link flags** field provides information about the physical link and displays one or more of the following values:

- **ACFC**—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.
- **Give-Up**—Link protocol does not continue connection attempts after repeated failures.
- **Loose-LCP**—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.
- **Loose-LMI**—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.
- **Loose-NCP**—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.
- **No-Keepalives**—Link protocol keepalives are disabled.
- **PFC**—Protocol field compression is configured. The PPP session negotiates the PFC option.

Logical Interface Flags Field

The **Logical interface flags** field provides information about the logical interface and displays one or more of the following values:

- **ACFC Encapsulation**—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).
- **Device-down**—Device has been administratively disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Clear-DF-Bit**—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.
- **Hardware-Down**—Interface protocol initialization failed to complete successfully.
- **PFC**—Protocol field compression is enabled for the PPP session.
- **Point-To-Point**—Interface is point-to-point.
- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Label-Switched Interface Traffic Statistics Field

When you use the **vrf-table-label** statement to configure a VRF routing table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table.

Any routes present in a VRF routing table and configured with the **vrf-table-label** statement are advertised with the LSI logical interface label allocated for the VRF routing table. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table. For more information on the **vrf-table-label** statement, including a list of supported interfaces, see the *Junos VPNs Configuration Guide*.

If you configure the **family mpls** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level and you also configure the **vrf-table-label** statement at the **[edit routing-instances routing-instance-name]** hierarchy level, the output for the **show interface interface-name extensive** command includes the following output fields about the LSI traffic statistics:

- **Input bytes**—Number of bytes entering the LSI and the current throughput rate in bits per second (bps).
- **Input packets**—Number of packets entering the LSI and the current throughput rate in packets per second (pps).

NOTE: If LSI interfaces are used with VPLS when **no-tunnel-services** is configured or L3VPN when **vrf-table-label** configuration is applied inside the routing-instance, the **Input packets** field associated with the core-facing interfaces may not display the correct value. Only the Input counter is affected because the LSI is used to receive traffic from the remote PEs. Traffic that arrives on an LSI interface might not be counted at both the Traffic Statistics and the Label-switched interface (LSI) traffic statistics levels.

This note applies to the following platforms:

- M Series routers with -E3 FPC model numbers or configured with an Enhanced CFEB (CFEB-E), and M120 routers
- MX Series routers with DPC or ADPC only

The following example shows the LSI traffic statistics that you might see as part of the output of the **show interface interface-name extensive** command:

Label-switched interface (LSI) traffic statistics:

Input bytes:	0	0 bps
Input packets:	0	0 pps

Policer Field

For the logical interface, the **Policer** field provides the policers that are to be evaluated when packets are received or transmitted on the interface. The format is **Policer: Input: *type-fpc/picport-in-policer*, Output: *type-fpc/pic/port-out-policer***. For example:

```
Policer: Input: at-1/2/0-in-policer, Output: at-2/4/0-out-policer
```

Protocol Field

For the logical interface, the **Protocol** field indicates the protocol family or families that are configured on the interface, displaying one or more of the following values:

- **aenet**—Aggregated Ethernet. Displayed on Fast Ethernet interfaces that are part of an aggregated Ethernet bundle.
- **ccc**—Circuit cross-connect (CCC). Configured on the logical interface of CCC physical interfaces.
- **inet**—IP version 4 (IPv4). Configured on the logical interface for IPv4 protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—IP version 6 (IPv6). Configured on the logical interface for IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), and BGP.
- **iso**—International Organization for Standardization (ISO). Configured on the logical interface for IS-IS traffic.
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI). Configured on the logical interface for link services bundling.
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end. Configured on the logical interface for multilink bundling.
- **mlppp**—Multilink Point-to-Point Protocol (MLPPP). Configured on the logical interface for multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). Configured on the logical interface for participation in an MPLS path.
- **pppoe**—Point-to-Point Protocol over Ethernet (PPPoE). Configured on Ethernet interfaces enabled to support multiple protocol families.
- **tcc**—Translational cross-connect (TCC). Configured on the logical interface of TCC physical interfaces.

- **tnp**—Trivial Network Protocol (TNP). 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.
- **vpls**—Virtual private LAN service (VPLS). Configured on the logical interface on which you configure VPLS.

RPF Failures Field

For the logical interface, the **RPF Failures** field provides information about the amount of incoming traffic (in packets and bytes) that failed a unicast reverse path forwarding (RPF) check on a particular interface. The format is **RPF Failures: Packets: xx,Bytes: yy**. For example:

RPF Failures: Packets: 0, Bytes:0

Source Class Field

For the logical interface, the **Source class** field provides the names of source class usage (SCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

Source class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980
(889)	(597762)
bronze	0	0
(0)	(0)
silver	0	0
(0)	(0)

clear interfaces statistics

Syntax

```
clear interfaces statistics (all | interface-name)
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 19.2R1 for QSFP-100GE-DWDM2 transceiver on MX10003, MX10008, MX10016, and MX204 routers.

Description

Set interface statistics to zero. If you issue the **clear interfaces statistics *interface-name*** command and then perform a graceful Routing Engine switchover, the interface statistics are not cleared on the new master. Reissue the command to clear the interface statistics again.

Starting in Junos OS Release 17.3R1, this command supports the clearing of Packet Forwarding Engine accounting statistics on logical interfaces configured with accounting options. On these interfaces, the current statistics values are stored as the new current baseline values and then the counters are reset to zero. If the **allow-clear** statement is included in the interface profile, then the cleared statistics values are reported to the accounting options flat file associated with the interface. Reporting is disabled by default; if **allow-clear** is not configured, then the CLI displays cleared statistics counters, but they are not reported to the flat file.

Starting in Junos OS Release 19.1R1, this command supports the clearing of unicast Reverse Path Forwarding (RPF) statistics.

Options

all—Set statistics on all interfaces to zero.

interface-name—Set statistics on a particular interface to zero.

Required Privilege Level

clear

List of Sample Output

[clear interfaces statistics on page 789](#)

Output Fields

When you enter this command, you are provided no feedback on the status of your request.

Sample Output

```
clear interfaces statistics
```

```
user@host> clear interfaces statistics
```

monitor interface

Syntax

```
monitor interface
<interface-name | traffic <detail>>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.

NOTE: On Junos OS Evolved, you can use the **monitor interface** command over SSH sessions, but console and Telnet sessions are not supported.

NOTE: This command is not supported on the QFX3000 QFabric switch.

Options

none—Display real-time statistics for all interfaces.

detail—(Optional) With traffic option only, display detailed output.

interface-name—(Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

traffic—(Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

Additional Information

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the c key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type

in the [CLI Explorer](#). To control the output of the **monitor interface** command while it is running, use the keys listed in [Table 70 on page 791](#). The keys are not case-sensitive.

Table 70: Output Control Keys for the monitor interface interface-name Command

Key	Action
c	Clears (returns to zero) the delta counters since monitor interface was started. This does not clear the accumulative counter. To clear the accumulative counter, use the clear interfaces interval command.
f	Freezes the display, halting the display of updated statistics and delta counters.
i	Displays information about a different interface. The command prompts you for the name of a specific interface.
n	Displays information about the next interface. The monitor interface command displays the physical or logical interfaces in the same order as the show interfaces terse command.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 71 on page 791](#). The keys are not case-sensitive.

Table 71: Output Control Keys for the monitor interface traffic Command

Key	Action
b	Displays the statistics in units of bytes and bits per second (bps).
c	Clears (return to 0) the delta counters in the Current Delta column. The statistics counters are not cleared.
d	Displays the Current Delta column (instead of the rate column) in bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or Esc	Quits the command and returns to the command prompt.
r	Displays the rate column (instead of the Current Delta column) in bps and pps.

Required Privilege Level

trace

List of Sample Output

[monitor interface \(Physical\) on page 793](#)

[monitor interface \(OTN Interface\) on page 795](#)

[monitor interface \(MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface\) on page 796](#)

[monitor interface \(MX480 Router with MPC5E and 100-Gigabit Ethernet Interface\) on page 797](#)

[monitor interface \(MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 798](#)

[monitor interface \(MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface\) on page 800](#)

[monitor interface \(MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 800](#)

[monitor interface \(Logical\) on page 801](#)

[monitor interface \(QFX3500 Switch\) on page 802](#)

[monitor interface traffic on page 803](#)

[monitor interface traffic \(QFX3500 Switch\) on page 803](#)

[monitor interface traffic detail \(QFX3500 Switch\) on page 804](#)

Output Fields

[Table 72 on page 792](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 72: monitor interface Output Fields

Field Name	Field Description	Level of Output
router1	Hostname of the router.	All levels
Seconds	How long the monitor interface command has been running or how long since you last cleared the counters.	All levels
Time	Current time (UTC).	All levels
Delay x/y/z	Time difference between when the statistics were displayed and the actual clock time. <ul style="list-style-type: none"> • x—Time taken for the last polling (in milliseconds). • y—Minimum time taken across all pollings (in milliseconds). • z—Maximum time taken across all pollings (in milliseconds). 	All levels
Interface	Short description of the interface, including its name, status, and encapsulation.	All levels
Link	State of the link: Up , Down , or Test .	All levels
Current delta	Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.	All levels

Table 72: monitor interface Output Fields (continued)

Field Name	Field Description	Level of Output
Local Statistics	<p>(Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	All levels
Remote Statistics	<p>(Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	All levels
Traffic statistics	<p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	All levels
Description	With the traffic option, displays the interface description configured at the [edit interfaces <i>interface-name</i>] hierarchy level.	detail

Sample Output

monitor interface (Physical)

```
user@host> monitor interface so-0/0/0
```

```

router1                               Seconds: 19                               Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC48

Traffic statistics:                               Current Delta
  Input packets:                               6045 (0 pps)                               [11]
  Input bytes:                               6290065 (0 bps)                               [13882]
  Output packets:                               10376 (0 pps)                               [10]
  Output bytes:                               10365540 (0 bps)                               [9418]

Encapsulation statistics:
  Input keepalives:                               1901                               [2]
  Output keepalives:                               1901                               [2]
  NCP state: Opened
  LCP state: Opened

Error statistics:
  Input errors:                               0                               [0]
  Input drops:                               0                               [0]
  Input framing errors:                               0                               [0]
  Policed discards:                               0                               [0]
  L3 incompletes:                               0                               [0]
  L2 channel errors:                               0                               [0]
  L2 mismatch timeouts:                               0                               [0]
  Carrier transitions:                               1                               [0]
  Output errors:                               0                               [0]
  Output drops:                               0                               [0]
  Aged packets:                               0                               [0]

Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                               1                               [0]
  LOF count                               1                               [0]
  SEF count                               1                               [0]
  ES-S                               0                               [0]
  SES-S                               0                               [0]

SONET statistics:
  BIP-B1                               458871                               [0]
  BIP-B2                               460072                               [0]
  REI-L                               465610                               [0]
  BIP-B3                               458978                               [0]
  REI-P                               458773                               [0]

Received SONET overhead:
  F1      : 0x00  J0      : 0x00  K1      : 0x00
  K2      : 0x00  S1      : 0x00  C2      : 0x00
  C2(cmp) : 0x00  F2      : 0x00  Z3      : 0x00

```

```

    Z4      : 0x00  S1(cmp)   : 0x00
Transmitted SONET overhead:
    F1      : 0x00  J0        : 0x01  K1        : 0x00
    K2      : 0x00  S1        : 0x00  C2        : 0xcf
    F2      : 0x00  Z3        : 0x00  Z4        : 0x00

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

monitor interface (OTN Interface)

user@host> **monitor interface ge-7/0/0**

```

Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:                0 (0 bps)
  Output bytes:               0 (0 bps)
  Input packets:              0 (0 pps)
  Output packets:             0 (0 pps)
Error statistics:
  Input errors:                0
  Input drops:                 0
  Input framing errors:        0
  Policed discards:            0
  L3 incompletes:              0
  L2 channel errors:           0
  L2 mismatch timeouts:        0
  Carrier transitions:         5
  Output errors:               0
  Output drops:                0
  Aged packets:                0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets              0
  Broadcast packets            0
  Multicast packets            0
  Oversized frames              0
  Packet reject count           0
  DA rejects                   0
  SA rejects                    0
Output MAC/Filter Statistics:

```

```

Unicast packets          0
Broadcast packets        0
Multicast packets        0
Packet pad count         0
Packet error count       0
OTN Link 0
  OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
  OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
  OTN OC - Seconds
    LOS                   2
    LOF                   9
  OTN OTU - FEC Statistics
    Corr err ratio        N/A
    Corr bytes            0
    Uncorr words          0
  OTN OTU - Counters
    BIP                   0
    BBE                   0
    ES                    0
    SES                   0
    UAS                   422
  OTN ODU - Counters
    BIP                   0
    BBE                   0
    ES                    0
    SES                   0
    UAS                   422
  OTN ODU - Received Overhead  APSPCC 0-3: 0

```

monitor interface (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

user@host> **monitor interface xe-0/0/3**

```

Interface: xe-0/0/3, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:          0 (0 bps)          [0]
  Output bytes:         0 (0 bps)          [0]
  Input packets:        0 (0 pps)          [0]
  Output packets:       0 (0 pps)          [0]
Error statistics:
  Input errors:         0                  [0]
  Input drops:          0                  [0]
  Input framing errors: 0                  [0]

```

```

    Policed discards:                0                [0]
    L3 incompletes:                  0                [0]
    L2 channel errors:               0                [0]
    L2 mismatch timeouts:            0                [0]
    Carrier transitions:              5                [0]
    Output errors:                   0                [0]
    Output drops:                    0                [0]
    Aged packets:                    0                [0]
Active alarms : None
Active defects: None
PCS statistics:                      Seconds
    Bit Errors                       0                [0]
    Errored blocks                   4                [0]
Input MAC/Filter statistics:
    Unicast packets                  0                [0]
    Broadcast packets                0                [0]
    Multicast packets                0                [0]
    Oversized frames                 0                [0]
    Packet reject count              0                [0]
    DA rejects                       0                [0]
    SA rejects                       0                [0]
Output MAC/Filter Statistics:
    Unicast packets                  0                [0]
    Broadcast packets                0                [0]
    Multicast packets                0                [0]
    Packet pad count                 0                [0]
    Packet error count               0                [0]

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface)

user@host> **monitor interface et-2/1/0**

```

Interface: et-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
Traffic statistics:                      Current delta
    Input bytes:                      0 (0 bps)        [0]
    Output bytes:                     0 (0 bps)        [0]
    Input packets:                    0 (0 pps)        [0]
    Output packets:                   0 (0 pps)        [0]
Error statistics:

```

```

Input errors:                                0                                [0]
Input drops:                                0                                [0]
Input framing errors:                        0                                [0]
Policed discards:                           0                                [0]
L3 incompletes:                             0                                [0]
L2 channel errors:                          0                                [0]
L2 mismatch timeouts:                       0                                [0]
Carrier transitions:                         263                             [0]
Output errors:                              0                                [0]
Output drops:                               0                                [0]
Aged packets:                               0                                [0]
OTN Link 0
  OTN Alarms:
  OTN Defects:
  OTN OC - Seconds
    LOS                                     129                                [0]
    LOF                                     2                                [0]
  OTN OTU - FEC Statistics
    Corr err ratio                         <8E-5
    Corr bytes                             169828399453                    [0]
    Uncorr words                           28939961456                    [0]
  OTN OTU - Counters
    BIP                                     0                                [0]
    BBE                                     0                                [0]
    ES                                     24                                [0]
    SES                                    0                                [0]
    UAS                                    1255                             [0]
  OTN ODU - Counters
    BIP                                     0                                [0]
    BBE                                     0                                [0]
    ES                                     24                                [0]
    SES                                    0                                [0]
    UAS                                    1256                             [0]
  OTN ODU - Received Overhead
    APSPCC 0-3:                           00 00 00 00                    [0]

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

monitor interface (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-6/1/0
```

```

Interface: xe-6/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
Input bytes: 0 (0 bps) [0]
Output bytes: 0 (0 bps) [0]
Input packets: 0 (0 pps) [0]
Output packets: 0 (0 pps) [0]
Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
Input framing errors: 0 [0]
Policed discards: 0 [0]
L3 incompletes: 0 [0]
L2 channel errors: 0 [0]
L2 mismatch timeouts: 0 [0]
Carrier transitions: 1 [0]
Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]
Active alarms : None
Active defects: None
PCS statistics:
Bit Errors 0 [0]
Errored blocks 1 [0]
Input MAC/Filter statistics:
Unicast packets 0 [0]
Broadcast packets 0 [0]
Multicast packets 0 [0]
Oversized frames 0 [0]
Packet reject count 0 [0]
DA rejects 0 [0]
SA rejects 0 [0]
Output MAC/Filter Statistics:
Unicast packets 0 [0]
Broadcast packets 0 [0]
Multicast packets 0 [0]
Packet pad count 0 [0]
Packet error count 0 [0]

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface et-9/0/0
```

```
Interface: et-9/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps

Traffic statistics:
Input bytes: 0 (0 bps) [0]
Output bytes: 0 (0 bps) [0]
Input packets: 0 (0 pps) [0]
Output packets: 0 (0 pps) [0]

Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
Input framing errors: 0 [0]
Policed discards: 0 [0]
L3 incompletes: 0 [0]
L2 channel errors: 0 [0]
L2 mismatch timeouts: 0 [0]
Carrier transitions: 1 [0]
Output errors: 0 [0]
Output drops: 0 [0]
Aged packets: 0 [0]
```

```
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-3/0/0
```

```
host name                      Seconds: 67                      Time: 23:46:46
                                Delay: 0/0/13

Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps

Traffic statistics:
Input bytes: 0 (0 bps) [0]
Output bytes: 0 (0 bps) [0]
Input packets: 0 (0 pps) [0]
Output packets: 0 (0 pps) [0]

Error statistics:
Input errors: 0 [0]
Input drops: 0 [0]
```



```

Input framing errors:          0          [0]
Policed discards:             0          [0]
L3 incompletes:               0          [0]
L2 channel errors:            0          [0]
L2 mismatch timeouts:         0          [0]
Carrier transitions:           3          [0]
Output errors:                 0          [0]
Output drops:                  0          [0]
Aged packets:                  0          [0]
OTN Link 0
  OTN Alarms:
  OTN Defects:
  OTN OC - Seconds
    LOS                        0          [0]
    LOF                        0          [0]
  OTN OTU - FEC Statistics
    Corr err ratio             N/A
    Corr bytes                  0          [0]
    Uncorr words                0          [0]
  OTN OTU - Counters
    BIP                        0
    BBE                        0          [0]
    ES                         0          [0]
    SES                        0          [0]
    UAS                        0          [0]
  OTN ODU - Counters
    BIP                        0
    BBE                        0          [0]
    ES                         0          [0]
    SES                        0          [0]
    UAS                        0          [0]
  OTN ODU - Received Overhead
    APSPCC 0-3:                00 00 00 00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

monitor interface (Logical)

user@host> **monitor interface so-1/0/0.0**

```

host name          Seconds: 16          Time: 15:33:39
                                          Delay: 0/0/1

Interface: so-1/0/0.0, Enabled, Link is Down

```

```

Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:
    Input bytes:          0          [0]
    Output bytes:         0          [0]
    Input packets:        0          [0]
    Output packets:       0          [0]
Remote statistics:
    Input bytes:          0 (0 bps)   [0]
    Output bytes:         0 (0 bps)   [0]
    Input packets:        0 (0 pps)   [0]
    Output packets:       0 (0 pps)   [0]
Traffic statistics:
    Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

monitor interface (QFX3500 Switch)

```
user@switch> monitor interface ge-0/0/0
```

```

Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:
    Input bytes:          0 (0 bps)   [0]
    Output bytes:         0 (0 bps)   [0]
    Input packets:        0 (0 pps)   [0]
    Output packets:       0 (0 pps)   [0]
Error statistics:
    Input errors:         0          [0]
    Input drops:          0          [0]
    Input framing errors: 0          [0]
    Policed discards:     0          [0]
    L3 incompletes:       0          [0]
    L2 channel errors:    0          [0]
    L2 mismatch timeouts: 0          [0]
    Carrier transitions:  0          [0]
    Output errors:        0          [0]
    Output drops:         0          [0]
    Aged packets:         0          [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
    Unicast packets       0          [0]

```

```
Broadcast packets          0 Multicast packet          [0]
```

```
Interface warnings:
```

```
o Outstanding LINK alarm
```

monitor interface traffic

```
user@host> monitor interface traffic
```

```
host name          Seconds: 15          Time: 12:31:09
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
so-1/0/0	Down	0	(0)	0	(0)
so-1/1/0	Down	0	(0)	0	(0)
so-1/1/1	Down	0	(0)	0	(0)
so-1/1/2	Down	0	(0)	0	(0)
so-1/1/3	Down	0	(0)	0	(0)
t3-1/2/0	Down	0	(0)	0	(0)
t3-1/2/1	Down	0	(0)	0	(0)
t3-1/2/2	Down	0	(0)	0	(0)
t3-1/2/3	Down	0	(0)	0	(0)
so-2/0/0	Up	211035	(1)	36778	(0)
so-2/0/1	Up	192753	(1)	36782	(0)
so-2/0/2	Up	211020	(1)	36779	(0)
so-2/0/3	Up	211029	(1)	36776	(0)
so-2/1/0	Up	189378	(1)	36349	(0)
so-2/1/1	Down	0	(0)	18747	(0)
so-2/1/2	Down	0	(0)	16078	(0)
so-2/1/3	Up	0	(0)	80338	(0)
at-2/3/0	Up	0	(0)	0	(0)
at-2/3/1	Down	0	(0)	0	(0)

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

monitor interface traffic (QFX3500 Switch)

```
user@switch> monitor interface traffic
```

```
switch          Seconds: 7          Time: 16:04:37
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
ge-0/0/0	Down	0	(0)	0	(0)
ge-0/0/1	Up	392187	(0)	392170	(0)

ge-0/0/2	Down	0	(0)	0	(0)
ge-0/0/3	Down	0	(0)	0	(0)
ge-0/0/4	Down	0	(0)	0	(0)
ge-0/0/5	Down	0	(0)	0	(0)
ge-0/0/6	Down	0	(0)	0	(0)
ge-0/0/7	Down	0	(0)	0	(0)
ge-0/0/8	Down	0	(0)	0	(0)
ge-0/0/9	Up	392184	(0)	392171	(0)
ge-0/0/10	Down	0	(0)	0	(0)
ge-0/0/11	Down	0	(0)	0	(0)
ge-0/0/12	Down	0	(0)	0	(0)
ge-0/0/13	Down	0	(0)	0	(0)
ge-0/0/14	Down	0	(0)	0	(0)
ge-0/0/15	Down	0	(0)	0	(0)
ge-0/0/16	Down	0	(0)	0	(0)
ge-0/0/17	Down	0	(0)	0	(0)
ge-0/0/18	Down	0	(0)	0	(0)
ge-0/0/19	Down	0	(0)	0	(0)
ge-0/0/20	Down	0	(0)	0	(0)
ge-0/0/21	Down	0	(0)	0	(0)
ge-0/0/22	Up	392172	(0)	392187	(0)
ge-0/0/23	Up	392185	(0)	392173	(0)
vcp-0	Down	0		0	
vcp-1	Down	0		0	
ae0	Down	0	(0)	0	(0)
bme0	Up	0		1568706	

monitor interface traffic detail (QFX3500 Switch)

user@switch> monitor interface traffic detail

```
switch
Time: 16:03:02
Seconds: 74
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
Description					
ge-0/0/0	Down	0	(0)	0	(0)
ge-0/0/1	Up	392183	(0)	392166	(0)
ge-0/0/2	Down	0	(0)	0	(0)
ge-0/0/3	Down	0	(0)	0	(0)
ge-0/0/4	Down	0	(0)	0	(0)
ge-0/0/5	Down	0	(0)	0	(0)
ge-0/0/6	Down	0	(0)	0	(0)
ge-0/0/7	Down	0	(0)	0	(0)

ge-0/0/8	Down	0	(0)	0	(0)
ge-0/0/9	Up	392181	(0)	392168	(0)
ge-0/0/10	Down	0	(0)	0	(0)
ge-0/0/11	Down	0	(0)	0	(0)
ge-0/0/12	Down	0	(0)	0	(0)
ge-0/0/13	Down	0	(0)	0	(0)
ge-0/0/14	Down	0	(0)	0	(0)
ge-0/0/15	Down	0	(0)	0	(0)
ge-0/0/16	Down	0	(0)	0	(0)
ge-0/0/17	Down	0	(0)	0	(0)
ge-0/0/18	Down	0	(0)	0	(0)
ge-0/0/19	Down	0	(0)	0	(0)
ge-0/0/20	Down	0	(0)	0	(0)
ge-0/0/21	Down	0	(0)	0	(0)
ge-0/0/22	Up	392169	(0)	392184	(1)
ge-0/0/23	Up	392182	(0)	392170	(0)
vcp-0	Down	0		0	
vcp-1	Down	0		0	
ae0	Down	0	(0)	0	(0)
bme0	Up	0		1568693	

request diagnostics tdr

Syntax

```
request diagnostics tdr (abort | start) interface interface-name
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Start a time domain reflectometry (TDR) diagnostic test on the specified interface. This test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 interfaces on line cards.

NOTE: We recommend running the TDR test when there is no traffic on the interface under test.

You view the results of the TDR test with the [show diagnostics tdr](#) command.

Options

abort—Stop the TDR test currently in progress on the specified interface. No results are reported, and previous results, if any, are cleared.

interface-name—The name of the interface.

start—Start a TDR test on the specified interface.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[show diagnostics tdr](#) | 812

[Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\)](#) | 400

List of Sample Output

[request diagnostics tdr start interface ge-0/0/19 on page 807](#)

Output Fields

[Table 73 on page 807](#) lists the output fields for the **request diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 73: request diagnostics tdr Output Fields

Field Name	Field Description
Test Status	<p>Information about the status of the TDR test request:</p> <ul style="list-style-type: none"> • Admin Down <i>interface-name</i>—The interface is administratively down. The TDR test cannot run on interfaces that are administratively down. • Interface <i>interface-name</i> not found—The interface does not exist. • Test successfully executed <i>interface-name</i>—The test has successfully started on the interface. You can view the test results with the show diagnostics tdr command. • VCT not supported on <i>interface-name</i>—The TDR test is not supported on the interface.

Sample Output

request diagnostics tdr start interface ge-0/0/19

```
user@switch> request diagnostics tdr start interface ge-0/0/19
```

```
Interface TDR detail:
```

```
Test status                : Test successfully executed ge-0/0/19
```

request chassis system-mode

Syntax

```
request chassis system-mode  
<mode-2x100G>  
<default-mode>
```

Release Information

Command introduced in Junos OS Release 19.3R1 for EX4300-48MP switches.

Description

Set the operating mode of the ports to 100-Gigabit Ethernet mode on an EX4300-48MP switch. The **request chassis system-mode** command sets the ports to 100-Gigabit Ethernet mode by overriding the **set chassis fpc fpc-slot pic pic-slot port port-number speed speed value** command. This command restarts the DC packet forwarding engine (DCPFE), disables the virtual chassis ports (VCP) and brings up both ports of the uplink module to operate at 100Gbps speed. Use the **show chassis system-mode** command to verify configured system mode and the **request chassis system-mode default-mode** command to return to the default operating mode.

Options

mode-2x100G—PIC mode configuration is 2x100Gbps. This will restart PFE and disable VCP ports.

default-mode—Default-mode is the oversubscribed mode. This will restart PFE.

Required Privilege Level

view

RELATED DOCUMENTATION

Uplink Modules in EX4300 Switches

[Setting the Operating Mode on a 2-Port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 Uplink Module](#) | 68

List of Sample Output

[request chassis system-mode on page 809](#)

[request chassis system-mode \(EX4300-48MP\) on page 809](#)

Sample Output

request chassis system-mode

user@switch> **request chassis system-mode mode-2x100G (mode - 2x100G on EX4300-48MP)**

System-mode has changed. This will restart PFE and disable VCP ports.

request chassis system-mode (EX4300-48MP)

user@switch> **request chassis system-mode default-mode**

System-mode has changed. This will restart PFE.

Show chassis system-mode

Syntax

```
show chassis system-mode
```

Release Information

Command introduced in Junos OS Release 19.3R1 for EX4300-48MP Switches.

Description

Displays the current system mode configuration on the switch. For EX4300-48MP the command, displays the current systemmode configured, for either the default mode or the current operating speed, 2x100G mode, for the switch.

Required Privilege Level

view

RELATED DOCUMENTATION

[Setting the Operating Mode on a 2-Port 40-Gigabit Ethernet QSFP+/100-Gigabit Ethernet QSFP28 Uplink Module | 68](#)

List of Sample Output

[show chassis system-mode \(EX4300-48MP\) on page 810](#)

[show chassis system-mode \(mode - 2x100G on EX4300-48MP\) on page 811](#)

Output Fields

[Table 74 on page 810](#) lists the output fields for the **show chassis system-mode** command.

Table 74: Output Fields for show chassis system-mode

Field Name	Field Description
Current System-Mode Configuration	Existing system operating mode for the device. For EX4300-48MP, the device can be in the default mode, or the 2x100G operating mode.

Sample Output

show chassis system-mode (EX4300-48MP)

user@switch> show chassis system-mode

```
localre:
```

```
-----
```

```
Current System-Mode Configuration: Default-mode
```

show chassis system-mode (mode - 2x100G on EX4300-48MP)

```
user@switch> show chassis system-mode
```

```
localre:
```

```
-----
```

```
Current System-Mode Configuration: 2x100G-mode
```

show diagnostics tdr

Syntax

```
show diagnostics tdr  
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the results of a time domain reflectometry (TDR) diagnostic test run on an interface. A TDR test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches— RJ-45 interfaces on line cards.

Use the [request diagnostics tdr](#) command to request a TDR test on a specified interface. Use the **show diagnostic tdr** command to display the last TDR test results for a specified interface or the last TDR test results for all network interfaces on the switch that support the TDR test.

Options

none—Show summarized last results for all interfaces on the switch that support the TDR test.

interface *interface-name*—(Optional) Show detailed last results for the specified interface or a range of interfaces. Specify a range of interfaces by entering the beginning and ending interface in the range, separated by a dash—for example, **ge-0/0/15-ge-0/0/20**.

Required Privilege Level

view

RELATED DOCUMENTATION

[request diagnostics tdr](#) | 806

[Diagnosing a Faulty Twisted-Pair Cable \(CLI Procedure\)](#) | 400

List of Sample Output

[show diagnostics tdr interface ge-0/0/19 \(Normal Cable\)](#) on page 815

[show diagnostics tdr interface ge-2/0/2 \(Faulty Cable\) on page 816](#)

[show diagnostics tdr \(All Supported Interfaces\) on page 817](#)

Output Fields

[Table 75 on page 813](#) lists the output fields for the **show diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 75: show diagnostics tdr Output Fields

Field Name	Field Description
Interface name or Interface	Name of interface for which TDR test results are being reported.
Test status	Status of TDR test: <ul style="list-style-type: none"> • Aborted—Test was terminated by operator before it was complete. • Failed—Test was not completed successfully. • Interface <i>interface-name</i> not found—Specified interface does not exist. • Not Started—No TDR test results are available for the interface. • Passed—Test completed successfully. The cable, however, might still have a fault—see the Cable status field for information on the cable. • Started—Test is currently running and not yet complete. • VCT not supported on <i>interface-name</i>—TDR test is not supported on the interface.
Link status	Operating status of link: UP or Down .
MDI pair	Twisted pair for which test results are being reported, identified by pin numbers. (Displayed only when the interface option is used.)

Table 75: show diagnostics tdr Output Fields (*continued*)

Field Name	Field Description
Cable status	<p>When detailed information is displayed, status for a twisted pair:</p> <ul style="list-style-type: none"> • Failed—TDR test failed on the cable pair. • Impedance Mismatch—Impedance on the twisted pair is not correct. Possible reasons for an impedance mismatch include: <ul style="list-style-type: none"> • The twisted pair is not connected properly. • The twisted pair is damaged. • The connector is faulty. • Normal—No cable fault detected for the twisted pair. • Open—Lack of continuity between the pins at each end of the twisted-pair. • Short on Pair-<i>n</i>—A short-circuit was detected on the twisted pair. <p>When summary information for all interfaces is displayed, status for the cable as a whole:</p> <ul style="list-style-type: none"> • Fault—A fault was detected on one or more of the twisted-pairs. • OK—No fault was detected on any of the twisted pairs.
Distance fault or	Distance to the fault in whole meters. If there is no fault, this value is 0.
Max distance fault	When summary information for all interfaces is displayed, this value is the distance to the most distant fault if there is more than one twisted pair with a fault.
Polarity swap	<p>Indicates the polarity status of the twisted pair:</p> <ul style="list-style-type: none"> • Normal—Polarity is normal. Each conductor in the twisted pair has been connected the same pins at the both ends of the connection. For example, a conductor connected to pin 1 at the near end of the connection is connected to pin 1 at the far end. • Reversed—Polarity has been reversed. For the twisted pair, the conductors have switched which pins they are connected to at the near and far ends of the connection. For example, the conductor connected to pin 1 at the near end is connected to pin 2 at the far end. <p>(Not available on EX8200 switches.) (Displayed only when the interface option is used)</p>
Skew time	Difference in nanoseconds between the propagation delay on this twisted pair and the twisted pair with the shortest propagation delay. (Not available on EX8200 switches.) (Displayed only when the interface option is used.)
Channel Pair	Number of the 10/100BASE-T transmit/receive pair being reported on.

Table 75: show diagnostics tdr Output Fields (*continued*)

Field Name	Field Description
Pair Swap	<p>Indicates whether or not the twisted pairs are swapped:</p> <ul style="list-style-type: none"> • MDI—The pairs are not swapped (straight-through cable). • MDIX—The pairs are swapped (cross-over cable). <p>(Displayed only when the interface option is used.)</p>
Downshift	<p>Indicates whether the connection speed is being downshifted:</p> <ul style="list-style-type: none"> • No Downshift—No downshifting of connection speed. • Downshift occurs—Connection speed is downshifted to 10 or 100 Mbs. This occurs if the cable is a two-pair cable rather than the four-pair cable required by Gigabit Ethernet. <p>(Displayed only when the interface option is used.)</p>

Sample Output

show diagnostics tdr interface ge-0/0/19 (Normal Cable)

```

user@switch> show diagnostics tdr interface ge-0/0/19
Interface TDR detail:
Interface name           : ge-0/0/19
Test status              : Passed
Link status              : UP
MDI pair                 : 1-2
  Cable status           : Normal
  Distance fault         : 0 Meters
  Polartiy swap          : Normal
  Skew time              : 0 ns
MDI pair                 : 3-6
  Cable status           : Normal
  Distance fault         : 0 Meters
  Polartiy swap          : Normal
  Skew time              : 8 ns
MDI pair                 : 4-5
  Cable status           : Normal
  Distance fault         : 0 Meters
  Polartiy swap          : Normal
  Skew time              : 8 ns

```

```

MDI pair           : 7-8
  Cable status      : Normal
  Distance fault    : 0 Meters
  Polartiy swap     : Normal
  Skew time         : 8 ns
Channel pair       : 1
  Pair swap         : MDI
Channel pair       : 2
  Pair swap         : MDI
Downshift          : No Downshift

```

show diagnostics tdr interface ge-2/0/2 (Faulty Cable)

```

user@switch> show diagnostics tdr interface ge-2/0/2
Interface TDR detail:
Interface name      : ge-2/0/2
Test status         : Passed
Link status         : Down
MDI Pair           : 1-2
  Cable status      : 1-2
  Distance fault    : 2 Meters
  Polartiy swap     : N/A
  Skew time         : N/A
MDI Pair           : 3-6
  Cable status      : Impedance Mismatch
  Distance fault    : 3 Meters
  Polartiy swap     : N/A
  Skew time         : N/A
MDI Pair           : 4-5
  Cable status      : Impedance Mismatch
  Distance fault    : 3 Meters
  Polartiy swap     : N/A
  Skew time         : N/A
MDI Pair           : 7-8
  Cable status      : Short on Pair-2
  Distance fault    : 3 Meters
  Polartiy swap     : N/A
  Skew time         : N/A
Channel pair       : 1
  Pair swap         : N/A
Channel pair       : 2
  Pair swap         : N/A
Downshift          : N/A

```


show diagnostics tdr (All Supported Interfaces)

```
user@switch> show diagnostics tdr
```

Interface	Test status	Link status	Cable status	Max distance	fault
ge-0/0/0	Not Started	N/A	N/A	N/A	
ge-0/0/1	Not Started	N/A	N/A	N/A	
ge-0/0/2	Started	N/A	N/A	N/A	
ge-0/0/3	Started	N/A	N/A	N/A	
ge-0/0/4	Passed	UP	OK	0	
ge-0/0/5	Passed	UP	Fault	173	
ge-0/0/6	Passed	UP	OK	0	
ge-0/0/7	Passed	UP	OK	0	
ge-0/0/8	Passed	UP	OK	0	
ge-0/0/9	Passed	UP	OK	0	
ge-0/0/10	Passed	UP	OK	0	
ge-0/0/11	Passed	UP	OK	0	
ge-0/0/12	Passed	UP	OK	0	
ge-0/0/13	Passed	UP	OK	0	
ge-0/0/14	Passed	UP	OK	0	
ge-0/0/15	Passed	UP	OK	0	
ge-0/0/16	Passed	UP	OK	0	
ge-0/0/17	Passed	UP	OK	0	
ge-0/0/18	Passed	UP	OK	0	
ge-0/0/19	Passed	UP	OK	0	
ge-0/0/20	Passed	Down	Fault	0	
ge-0/0/21	Passed	Down	Fault	5	
ge-0/0/22	Passed	UP	OK	0	
ge-0/0/23	Passed	UP	OK	0	

show forwarding-options enhanced-hash-key

Syntax

```
show forwarding-options enhanced-hash-key
```

Release Information

Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.

Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

Fabric Load Balancing Options output fields introduced in Junos OS Release 14.1X53-D10.

Incoming port output field introduced in Junos OS Release 18.4R1 for QFX10000 Series switches.

Command introduced in Junos OS Release 19.4R1 for QFX5120-32C and QFX5120-48Y switches.

The **ecmp-dlb** statement introduced in Junos OS evolved Release 19.4R2 for QFX5220 switches.

Description

Display information about which packet fields are used by the hashing algorithm to make hashing decisions.

You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the **forwarding-options enhanced-hash-key** statement.

Required Privilege Level

view

RELATED DOCUMENTATION

[Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) | 349](#)

[Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic | 342](#)

List of Sample Output

[show forwarding-options enhanced-hash-key \(Layer 2 Payload Hash Mode\) on page 821](#)

[show forwarding-options enhanced-hash-key \(Layer 2 Header Hash Mode\) on page 822](#)

[show forwarding-options enhanced-hash-key \(Fabric Load Balancing Options\) on page 822](#)

[show forwarding-options enhanced-hash-key \(Dynamic Load Balancing Options\) on page 822](#)

[show forwarding-options enhanced-hash-key \(QFX10000 Series Switches\) on page 823](#)

Output Fields

[Table 76 on page 819](#) lists the output fields for the **show forwarding-options enhanced-hash-key** command. Output fields are listed in the approximate order in which they first appear. Output fields vary by platform.

Table 76: show forwarding-options enhanced-hash-key Output Fields

Field Name	Field Description
Hash-Mode	Current hash mode: Layer 2 header or Layer 2 payload.
Protocol	Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.
Destination L4 Port	Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No.
Source L4 Port	Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.
Destination IPv4 Addr	Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
Source IPv4 Addr	Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
Incoming port	Indicates whether the incoming port number (interface) is or is not used by the hashing algorithm. Yes or No NOTE: When passive monitoring is enabled on a QFX10000 Series switch interface, the inet , inet6 and L2 fields are all set to No.
Vlan id	Indicates whether the Vlan ID field is or is not used by the hashing algorithm: Yes or No.
Inner-Vlan ID	Indicates whether the inner Vlan field is or is not used by the hashing algorithm: Yes or No.
Next Hdr	Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.
Destination IPv6 Addr	Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
Source IPv6 Addr	Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
Ether Type	Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.

Table 76: show forwarding-options enhanced-hash-key Output Fields (*continued*)

Field Name	Field Description
Destination MAC Address	Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.
Source MAC Address	Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.
Load Balancing Method for ALB	<p>Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet.</p> <p>The load balancing method is flowlet by default, and can be configured using the fabric-load-balance statement.</p>
Load Balancing Method for DLB (QFX5120-32C, QFX5120-48Y, and QFX5220 switches)	<p>Indicates various Dynamic Load Balancing (DLB) modes:</p> <ul style="list-style-type: none"> • Flowlet • Assigned flow • Per-packet <p>Refer “Dynamic Load Balancing” on page 372 for more details.</p>
Fabric Link Scale	Indicates the fabric link scale, in mbps.
Inactivity Interval	<p>Indicates the fabric load balance inactivity interval, in microseconds (us).</p> <p>The inactivity interval is 16 microseconds by default, and can be configured using the inactivity-interval statement.</p>
Hash Region Size/Trunk	Indicates the hash region size, in buckets per fabric trunk.
Seed	A hash seed value, between 0 and 4294967295. If a hash-seed value is not configured it is automatically assigned on the QFX10000 Series switches. A hash-seed prevents traffic polarization to same links on the next hop QFX switch when two are connected with LAG/ECMP.
Key	Indicates whether the GRE key field is or is not used by the hashing algorithm: Yes or No.
Protocol	Indicates if a Generic Router Encapsulation (GRE) endpoint over routes was dynamically learned by a routing protocol such as RIP or OSPF.
MPLS Enabled	Indicates if MPLS is enabled under L2 switching.

Table 76: show forwarding-options enhanced-hash-key Output Fields (*continued*)

Field Name	Field Description
VXLAN VNID	A 24-bit virtual network identifier (VNID) that uniquely identifies the Virtual Extensible Local Area Networks (VXLAN) segment.

Sample Output

show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

user@switch> **show forwarding-options enhanced-hash-key**

Slot 0

Current Hash Settings

Hash-Mode :layer2-payload

inet Hash settings-

inet packet fields

Protocol	: Yes
Destination L4 Port	: Yes
Source L4 Port	: Yes
Destination IPv4 Addr	: Yes
Source IPv4 Addr	: Yes
Vlan id	: No

inet6 Hash settings-

inet6 packet fields

Next Hdr	: Yes
Destination L4 Port	: Yes
Source L4 Port	: Yes
Destination IPv6 Addr	: Yes
Source IPv6 Addr	: Yes
Vlan id	: No

show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
```

```
Slot 0

Current Hash Settings
-----

Hash-Mode                               : layer2-header

layer2 Hash settings-
-----
layer2 packet fields
  Ether Type                            : Yes
  Destination MAC Address               : Yes
  Source MAC Address                    : Yes
  VLAN ID                               : No
```

show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
```

```
<some output removed for brevity>

Fabric Load Balancing Options
-----

Load Balancing Method   : Flowlet
Fabric Link Scale        : 40960 (mbps)
Inactivity Interval     : 16 (us)
Hash Region Size/Trunk  : 1024 (buckets)
```

show forwarding-options enhanced-hash-key (Dynamic Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key ecmp-dlb
```

```
Slot 0
Current RTAG7 Settings
-----

Hash-Mode                               : layer2-payload
inet RTAG7 settings-
-----
```

```

inet packet fields
  Protocol                : Yes
  Destination L4 Port     : Yes
  Source L4 Port          : Yes
  Destination IPv4 Addr   : Yes
  Source IPv4 Addr        : Yes
  Vlan id                 : No

```

```
inet6 RTAG7 settings-
```

```
-----
```

```

inet6 packet fields
  Next Hdr                : Yes
  Destination L4 Port     : Yes
  Source L4 Port          : Yes
  Destination IPv6 Addr   : Yes
  Source IPv6 Addr        : Yes
  Vlan id                 : No

```

```
ECMP Load Balancing Options
```

```
-----
```

```

Load Balancing Method    : Flowlet
Inactivity Interval      : 64 (us)

```

show forwarding-options enhanced-hash-key (QFX10000 Series Switches)

```
user@switch> show forwarding-options enhanced-hash-key
```

```
Slot 0
```

```

Seed value for Hash function    0: 2301323130
Seed value for Hash function    1: 2301323130
Seed value for Hash function    2: 2301323130
Seed value for Hash function    3: 2301323130

```

```
Inet settings:
```

```
-----
```

```

  IPV4 dest address:    Yes
  IPV4 source address:  Yes
  L4 Dest Port:         Yes
  L4 Source Port:       Yes
  Incoming port:        No

```

```
Inet6 settings:
```

```
-----
```

```
        IPV6 dest address:    Yes
        IPV6 source address:  Yes
        L4 Dest Port:        Yes
        L4 Source Port:      Yes
        Incoming port:       No
L2 settings:
-----
        Dest Mac address:    No
        Source Mac address:  No
        Vlan Id:             Yes
        Inner-vlan Id:       No
        Incoming port:       No
GRE settings:
-----
        Key:                 No
        Protocol:            No
MPLS settings:
-----
MPLS Enabled:                Yes

VXLAN settings:
-----
        VXLAN VNID:         No
```


show interfaces (Discard)

Syntax

```
show interfaces dsc
<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 discard interface.

Options

dsc—Display standard information about the specified discard interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—This option is not relevant for the discard interface and always shows a value of 0.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) This option is not relevant for the discard interface and always shows a value of 0.

Required Privilege Level

view

RELATED DOCUMENTATION

[show interfaces \(ATM\)](#)

[show interfaces routing](#)

List of Sample Output

[show interfaces dsc on page 829](#)

[show interfaces dsc brief on page 830](#)

[show interfaces dsc detail on page 830](#)

[show interfaces dsc extensive on page 831](#)

Output Fields

Table 77 on page 826 lists the output fields for the **show interfaces** (discard) command. Output fields are listed in the approximate order in which they appear.

Table 77: Discard show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, whether the interface is enabled, and the state of the physical interface: Up or Down .	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Type	Type of interface. Software-Pseudo indicates a standard software interface with no associated hardware device.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Clocking	Reference clock source. It can be Internal or External .	brief detail extensive
Speed	Speed at which the interface is running.	brief detail extensive
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 779 .	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under "Common Output Fields Description" on page 779 .	All levels
Link type	Encapsulation being used on the physical interface.	detail extensive
Link flags	Information about the link. Possible values are described in the "Link Flags" section under "Common Output Fields Description" on page 779 .	detail extensive
Physical info	Information about the physical interface.	detail extensive

Table 77: Discard show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	detail extensive
Current address, Hardware address	Configured MAC address and hardware MAC address.	detail extensive
Alternate link address	Backup address of the link.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive

Table 77: Discard show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface:</p> <ul style="list-style-type: none"> • Errors—Sum of incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive
Output errors	<p>(Extensive only) Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive

Table 77: Discard show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family configured on the logical interface, such as iso , inet6 , or mpls .	All levels
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive

Sample Output

show interfaces dsc

user@host> **show interfaces dsc**

```
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5
  Type: Software-Pseudo, MTU: Unlimited
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags       : None
  Last flapped    : Never
    Input packets : 0
    Output packets: 0

Logical interface dsc.0 (Index 66) (SNMP ifIndex 235)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  Protocol inet, MTU: Unlimited
    Flags: None
```

show interfaces dsc brief

```
user@host> show interfaces dsc brief
```

```
Physical interface: dsc, Enabled, Physical link is Up
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
Unspecified, Speed: Unspecified
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface dsc.0
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  inet
```

show interfaces dsc detail

```
user@host> show interfaces dsc detail
```

```
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5, Generation: 9
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
Unspecified, Speed: Unspecified
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type         : Unspecified
  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
    Output bytes : 0
    Input  packets: 0
    Output packets: 0

Logical interface dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
  Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0
  Flags: None
```

show interfaces dsc extensive

```
user@host> show interfaces dsc extensive
```

```
Physical interface: dsc, Enabled, Physical link is Up
  Interface index: 5, SNMP ifIndex: 5, Generation: 9
  Type: Software-Pseudo, Link-level type: Unspecified, MTU: Unlimited, Clocking:
Unspecified, Speed: Unspecified
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type         : Unspecified
  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
    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 dsc.0 (Index 66) (SNMP ifIndex 235) (Generation 6)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Unspecified
    Protocol inet, MTU: Unlimited, Generation: 14, Route table: 0
```

show interfaces

List of Syntax

[Syntax \(Gigabit Ethernet\) on page 832](#)

[Syntax \(10 Gigabit Ethernet\) on page 832](#)

[Syntax \(ACX5448, ACX5448-D, ACX710\) on page 832](#)

[Syntax \(SRX Series Devices and \(vSRX and vSRX 3.0 platforms\)\) on page 832](#)

Syntax (Gigabit Ethernet)

```
show interfaces ge-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (10 Gigabit Ethernet)

```
show interfaces xe-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (ACX5448, ACX5448-D, ACX710)

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (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.

SRX4600 supports 40-Gigabit Ethernet breakouts only in PIC mode. Use the **show interfaces extensive** command to view the speed configured for the interface on SRX4600. Reboot the device for the changed configuration to take effect.

On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Starting in Junos OS Release 18.4R1, Output fields **Next-hop** and **vpls-status** is displayed in the **show interfaces interface-name detail** command, only for Layer 2 protocols on MX480 routers.

In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the **show interfaces interface-name <media> <extensive>** command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed. In the earlier Junos Releases, incorrect autonegotiation status was displayed even when the autonegotiation was disabled.

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
19.2R3	In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the show interfaces <i>interface-name</i> <media> <extensive> command displays the autonegotiation status only for the interface that supports autonegotiation.
18.4R1	Starting in Junos OS Release 18.4R1, Output fields Next-hop and vpls-status is displayed in the show interfaces <i>interface name</i> detail command, only for Layer 2 protocols on MX480 routers.

RELATED DOCUMENTATION

Understanding Layer 2 Interfaces on Security Devices

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

List of Sample Output

[show interfaces terse \(ACX5448, ACX5448-D, ACX710 channelized interface\) on page 886](#)

[show interfaces \(Gigabit Ethernet\) on page 887](#)

[show interfaces \(Gigabit Ethernet on MX Series Routers\) on page 887](#)

[show interfaces \(link degrade status\) on page 888](#)

[show interfaces extensive \(Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration\) on page 889](#)

[show interfaces brief \(Gigabit Ethernet\) on page 890](#)
[show interfaces detail \(Gigabit Ethernet\) on page 891](#)
[show interfaces extensive \(Gigabit Ethernet IQ2\) on page 892](#)
[show interfaces \(Gigabit Ethernet Unnumbered Interface\) on page 896](#)
[show interfaces \(ACI Interface Set Configured\) on page 897](#)
[show interfaces \(ALI Interface Set\) on page 897](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, IQ2\) on page 898](#)
[show interfaces extensive \(10-Gigabit Ethernet, WAN PHY Mode\) on page 901](#)
[show interfaces extensive \(10-Gigabit Ethernet, DWDM OTN PIC\) on page 903](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode\) on page 907](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only\) on page 907](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only\) on page 908](#)
[Sample Output SRX Gigabit Ethernet on page 910](#)
[Sample Output SRX Gigabit Ethernet on page 911](#)
[show interfaces \(Gigabit Ethernet for vSRX and vSRX 3.0\) on page 911](#)
[show interfaces detail \(Gigabit Ethernet\) on page 912](#)
[show interfaces statistics st0.0 detail on page 914](#)
[show interfaces extensive \(Gigabit Ethernet\) on page 915](#)
[show interfaces terse on page 919](#)
[show interfaces terse \(vSRX and vSRX 3.0\) on page 920](#)
[show interfaces controller \(Channelized E1 IQ with Logical E1\) on page 920](#)
[show interfaces controller \(Channelized E1 IQ with Logical DS0\) on page 920](#)
[show interfaces descriptions on page 920](#)
[show interfaces destination-class all on page 921](#)
[show interfaces diagnostics optics on page 921](#)
[show interfaces far-end-interval coc12-5/2/0 on page 922](#)
[show interfaces far-end-interval coc1-5/2/1:1 on page 923](#)
[show interfaces filters on page 923](#)
[show interfaces flow-statistics \(Gigabit Ethernet\) on page 924](#)
[show interfaces interval \(Channelized OC12\) on page 925](#)
[show interfaces interval \(E3\) on page 926](#)
[show interfaces interval \(SONET/SDH\) \(SRX devices\) on page 926](#)
[show interfaces load-balancing \(SRX devices\) on page 927](#)
[show interfaces load-balancing detail \(SRX devices\) on page 927](#)
[show interfaces mac-database \(All MAC Addresses on a Port SRX devices\) on page 927](#)
[show interfaces mac-database \(All MAC Addresses on a Service SRX devices\) on page 928](#)
[show interfaces mac-database mac-address on page 929](#)
[show interfaces mc-ae \(SRX devices\) on page 929](#)
[show interfaces media \(SONET/SDH\) on page 930](#)
[show interfaces policers \(SRX devices\) on page 930](#)
[show interfaces policers interface-name \(SRX devices\) on page 931](#)

[show interfaces queue \(SRX devices\) on page 931](#)
[show interfaces redundancy \(SRX devices\) on page 932](#)
[show interfaces redundancy \(Aggregated Ethernet SRX devices\) on page 933](#)
[show interfaces redundancy detail \(SRX devices\) on page 933](#)
[show interfaces routing brief \(SRX devices\) on page 933](#)
[show interfaces routing detail \(SRX devices\) on page 934](#)
[show interfaces routing-instance all \(SRX devices\) on page 935](#)
[show interfaces snmp-index \(SRX devices\) on page 935](#)
[show interfaces source-class all \(SRX devices\) on page 935](#)
[show interfaces statistics \(Fast Ethernet SRX devices\) on page 936](#)
[show interfaces switch-port \(SRX devices\) on page 937](#)
[show interfaces transport pm \(SRX devices\) on page 937](#)
[show security zones \(SRX devices\) on page 939](#)

Output Fields

Table 78 on page 838 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 79 on page 877.

Table 78: show interfaces (Gigabit Ethernet) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “ Common Output Fields Description ” on page 779.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	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
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 779 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “Common Output Fields Description” on page 779 .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second:timezone (hour:minute:second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail 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
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
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. 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. • Output bytes—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. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <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 show interfaces command.</p>	detail extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (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

Table 78: 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"> • 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. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Drops 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"> • Collisions—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. • Aged packets—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. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Egress queues	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: 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 show interfaces 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>	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Dropped packets field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>	detail extensive
Ingress queues	<p>Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.</p>	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the 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 None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Interface transmit statistics	<p>(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—Count of corrected errors in the last second. • Corrected Error Ratio—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
PCS statistics	<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"> • Bit errors—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errorred blocks—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link Degrade	<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"> • Link Monitoring—Indicates if physical link degrade monitoring is enabled on the interface. <ul style="list-style-type: none"> • Enable—Indicates that link degrade monitoring has been enabled (using the link-degrade-monitor statement) on the interface. • Disable—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. • Link Degrade Set Threshold—The BER threshold value at which the link is considered degraded and a corrective action is triggered. • Link Degrade Clear Threshold—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted. • Estimated BER—The estimated bit error rate. • Link-degrade event—Shows link degrade event information. <ul style="list-style-type: none"> • Seconds—Time (in seconds) elapsed after a link degrade event occurred. • Count—The number of link degrade events recorded. • State—Shows the link degrade status (example: Defect Active). 	detail extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC statistics		extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces command. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds interface MTU, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. 	

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces 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"> • Code violations—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.” 	
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

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field 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. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field must not increment. • CAM destination filters, CAM source filters—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. 	extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
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"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 78: 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"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner—Information from the remote Ethernet device: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the link partner, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), Symmetric/Asymmetric (link partner supports PAUSE on receive and transmit or only PAUSE on transmit), and None (link partner does not support flow control). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the local Ethernet device: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are Symmetric/Asymmetric (local device supports PAUSE on receive and transmit or only PAUSE on receive) and None (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 Symmetric (local device supports PAUSE on receive and transmit), Asymmetric (local device supports PAUSE on receive), and None (local device does not support flow control). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Received path trace, Transmitted path trace	(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.	extensive
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—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. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—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. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux	<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"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels

Table 78: 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 agent-circuit-id autoconfiguration stanza.</p> <p>Dynamic Profile—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>NOTE: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the line-identity autoconfiguration stanza.</p>	brief detail extensive none
Line Identity	<p>Information displayed for access-line-identifier (ALI) interface sets configured with the line-identity autoconfiguration stanza.</p> <ul style="list-style-type: none"> • Dynamic Profile—Name of the dynamic profile that defines the ALI interface set. • Trusted option used to create the ALI interface set: Circuit-id, Remote-id, or Accept-no-ids. More than one option can be configured. <p>If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information.</p> <p>NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the agent-circuit-id autoconfiguration stanza.</p>	detail
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 779 .	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 inet6 under logical interface statistics.</p> <ul style="list-style-type: none"> • Max nh cache—Maximum interface neighbor discovery nexthop cache size. • New hold nh limit—Maximum number of new unresolved nexthops. • Curr nh cnt—Current number of resolved nexthops in the NDP queue. • Curr new hold cnt—Current number of unresolved nexthops in the NDP queue. • NH drop cnt—Number of NDP requests not serviced. 	All levels

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Dynamic Profile	Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	detail extensive none
Service Name Table	Name of the service name table for the interface configured with a PPPoE family.	detail extensive none
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none
Duplicate Protection	State of PPPoE duplicate protection: On or Off . 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.	detail extensive none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: On or Off . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	detail extensive none
AC Name	Name of the access concentrator.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—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. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: 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 Output bytes and Output packets 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>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 779 .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	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.	detail extensive
Output Filters	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.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none

Table 78: show interfaces (Gigabit Ethernet) Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

The following table describes the output fields for the **show interfaces** (10-Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 779 .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels

Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	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
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 779 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “Common Output Fields Description” on page 779 .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none

Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail 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 extensive
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	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"> • Input bytes—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. • Output bytes—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. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • 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.

Egress queues	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: 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 show interfaces 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>	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Ingress queues	<p>Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.</p>	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the 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 None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN alarms	<p>Active OTN alarms identified on the interface.</p>	detail extensive

OTN defects	OTN defects received on the interface.	detail extensive
OTN FEC Mode	<p>The FECmode configured on the interface.</p> <ul style="list-style-type: none"> • efec—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors. • gfec—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • none—FEC mode is not configured. 	detail extensive
OTN Rate	<p>OTN mode.</p> <ul style="list-style-type: none"> • fixed-stuff-bytes—Fixed stuff bytes 11.0957 Gbps. • no-fixed-stuff-bytes—No fixed stuff bytes 11.0491 Gbps. • pass-through—Enable OTN passthrough mode. • no-pass-through—Do not enable OTN passthrough mode. 	detail extensive
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: enabled or disabled .	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
OTN FEC alarms	<p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <ul style="list-style-type: none"> • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect. 	detail extensive
OTN OC	<p>OTN OC defects triggered on the interface.</p> <ul style="list-style-type: none"> • LOS—OC Loss of Signal defect. • LOF—OC Loss of Frame defect. • LOM—OC Loss of Multiframe defect. • Wavelength Lock—OC Wavelength Lock defect. 	detail extensive

OTN OTU	<p>OTN OTU defects detected on the interface</p> <ul style="list-style-type: none"> • AIS—OTN AIS alarm. • BDI—OTN OTU BDI alarm. • IAE—OTN OTU IAE alarm. • TTIM—OTN OTU TTIM alarm. • SF—OTN ODU bit error rate fault alarm. • SD—OTN ODU bit error rate defect alarm. • TCA-ES—OTN ODU ES threshold alarm. • TCA-SES—OTN ODU SES threshold alarm. • TCA-UAS—OTN ODU UAS threshold alarm. • TCA-BBE—OTN ODU BBE threshold alarm. • BIP—OTN ODU BIP threshold alarm. • BBE—OTN OTU BBE threshold alarm. • ES—OTN OTU ES threshold alarm. • SES—OTN OTU SES threshold alarm. • UAS—OTN OTU UAS threshold alarm. 	detail extensive
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	detail extensive
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	detail extensive
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	detail extensive
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	detail extensive
PCS statistics	<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"> • Bit errors—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errored blocks—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive

MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	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>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—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. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. 	extensive

WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<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"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	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">• Seconds—Number of seconds the defect has been active.• Count—Number of times that the defect has gone from inactive to active.• State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none">• BIP-B3—Bit interleaved parity for SONET section overhead• REI-P—Remote error indication• LOP-P—Loss of pointer (path)• AIS-P—Path alarm indication signal• RDI-P—Path remote defect indication• UNEQ-P—Path unequipped• PLM-P—Path payload label mismatch• ES-P—Errored seconds (near-end STS path)• SES-P—Severely errored seconds (near-end STS path)• UAS-P—Unavailable seconds (near-end STS path)• SES-PFE—Severely errored seconds (far-end STS path)• UAS-PFE—Unavailable seconds (far-end STS path)	extensive
----------	--	-----------

Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, 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

Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive

Logical Interface

Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels

VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—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. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—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. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<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"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 779 .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none

Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—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. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: 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 Output bytes and Output packets 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>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 779 .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	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.	detail extensive

Output Filters	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.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

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 79: 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	show interfaces ge-0/3/0 extensive	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: 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	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Table 80 on page 878 lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 80: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels

Table 80: 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 Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	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 None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
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"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 80: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface.</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (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. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 80: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface.</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; therefore, for Gigabit Ethernet PICs, this number must 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 must 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 interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters and queue number	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 80: 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"> • Total octets and total packets—Total number of octets and packets. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive

Table 80: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—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. 	extensive

Table 80: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. 	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Interface transmit statistics	<p>Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p>	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 80: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—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. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
Local statistics	Number and rate of bytes and packets destined to the device.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: 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 Output bytes and Output packets 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>	extensive
Security	Security zones that interface belongs to.	extensive
Flow Input statistics	Statistics on packets received by flow module.	extensive

Table 80: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flow Output statistics	Statistics on packets sent by flow module.	extensive
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	extensive
Protocol	Protocol family.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. .	detail extensive
Addresses, Flags	Information about the address flags..	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output Gigabit Ethernet

show interfaces terse (ACX5448, ACX5448-D, ACX710 channelized interface)

```
user@host> show interfaces terse et-0/1/2
```

Interface	Admin	Link	Proto	Local	Remote
et-0/1/2:0			up		down
et-0/1/2:1			up		down
et-0/1/2:2			up		down
et-0/1/2:3			up		down

show interfaces (Gigabit Ethernet)

user@host> **show interfaces ge-3/0/2**

```
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues       : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped    : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

user@host> **show interfaces ge-2/2/2**

```
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
```

```

Link flags      : None
CoS queues      : 8 supported, 4 maximum usable queues
Schedulers     : 0
Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects  : None
Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
        Flags: 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           Seconds
  Bit errors           0
  Errored blocks       0
Link Degrad* :
Link Monitoring           : Enable
Link Degrad Set Threshold: : 1E-7
Link Degrad 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:    35384446506          1420444 pps
Direction : Output
Interface transmit statistics: Enabled

Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
Output bytes      :          195560312716          522726272 bps
Output packets:    4251311146          1420451 pps

```

```
user@host> show interfaces ge-5/2/0.0 statistics detail
```

```

Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
Input  bytes      :          271524
Output bytes      :          37769598
Input  packets:    3664
Output packets:    885790
IPv6 transit statistics:

```

```

    Input  bytes :                0
    Output bytes :            16681118
    Input  packets:                0
    Output packets:            362633
Local statistics:
    Input  bytes :            271524
    Output bytes :            308560
    Input  packets:            3664
    Output packets:            3659
Transit statistics:
    Input  bytes :                0                0 bps
    Output bytes :            37461038            0 bps
    Input  packets:                0                0 pps
    Output packets:            882131            0 pps
IPv6 transit statistics:
    Input  bytes :                0                0 bps
    Output bytes :            16681118            0 bps
    Input  packets:                0                0 pps
    Output packets:            362633            0 pps

```

show interfaces brief (Gigabit Ethernet)

user@host> **show interfaces ge-3/0/2 brief**

```

Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface ge-3/0/2.0
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  ccc

Logical interface ge-3/0/2.32767
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

```

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
```

```
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : None
  CoS queues       : 4 supported, 4 maximum usable queues
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped    : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
    Input  bytes :                0                0 bps
    Input  packets:                0                0 pps
    Drop  bytes  :                0                0 bps
    Drop  packets:                0                0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   0                    0                    0
    1 expedited-fo  0                    0                    0
    2 assured-forw  0                    0                    0
    3 network-cont  0                    0                    0
  Egress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   0                    0                    0
    1 expedited-fo  0                    0                    0
    2 assured-forw  0                    0                    0
    3 network-cont  0                    0                    0
  Active alarms  : None
  Active defects : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
```

```

Out(swap-push 0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Local statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Transit statistics:
    Input  bytes   :                0                0 bps
    Output bytes   :                0                0 bps
    Input  packets :                0                0 pps
    Output packets :                0                0 pps
  Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
  Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
  Traffic statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Local statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Transit statistics:
    Input  bytes   :                0                0 bps
    Output bytes   :                0                0 bps
    Input  packets :                0                0 pps
    Output packets :                0                0 pps

```

show interfaces extensive (Gigabit Ethernet IQ2)

user@host> show interfaces ge-7/1/3 extensive

```

Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4004000
  Link flags       : None
  CoS queues       : 8 supported, 4 maximum usable queues
  Schedulers       : 256
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
  Last flapped    : 2007-11-07 21:31:41 PST (02:03:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes      :          38910844056          7952 bps
    Output bytes     :          7174605          8464 bps
    Input packets    :          418398473          11 pps
    Output packets   :          78903          12 pps
  IPv6 transit statistics:
    Input bytes      :          0
    Output bytes     :          0
    Input packets    :          0
    Output packets   :          0
  Ingress traffic statistics at Packet Forwarding Engine:
    Input bytes      :          38910799145          7952 bps
    Input packets    :          418397956          11 pps
    Drop bytes       :          0          0 bps
    Drop packets     :          0          0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort    418390823          418390823          0
    1 expedited-fo    0                  0                  0
    2 assured-forw    0                  0                  0
    3 network-cont    7133              7133              0
  Egress queues: 4 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

```

```

0 best-effort          1031          1031          0
1 expedited-fo         0            0            0
2 assured-forw         0            0            0
3 network-cont        77872         77872          0
Active alarms   : None
Active defects  : None
MAC statistics:
    Receive      Transmit
Total octets    38910844056    7174605
Total packets   418398473      78903
Unicast packets 408021893366      1026
Broadcast packets      10         12
Multicast packets 418398217      77865
CRC/Align errors      0            0
FIFO errors           0            0
MAC control frames    0            0
MAC pause frames      0            0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations       0  OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
Payload Type: 0x08
OTN Transmitted Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08
Filter statistics:
    Input packet count    418398473
    Input packet rejects   479
    Input DA rejects       479
    Input SA rejects       0
    Output packet count    78903
    Output packet pad count      0
    Output packet error count    0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
    Remote fault: OK
Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 7

```

CoS information:

Direction : Output

CoS transmit queue	Bandwidth	Buffer	Priority	Limit
	% bps	%	usec	
0 best-effort	95 950000000	95	0	low none
3 network-control	5 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 78 on page 838](#).

show interfaces (Gigabit Ethernet Unnumbered Interface)

user@host> **show interfaces ge-3/2/0**

```

Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : None
  CoS queues       : 8 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:f8, Hardware address: 00:00:5e:00:53:f8
  Last flapped    : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate       : 0 bps (0 pps)
  Output rate      : 624 bps (1 pps)
  Active alarms    : None
  Active defects   : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 6
  Protocol inet, MTU: 1500
  Flags: Unnumbered
  Donor interface: lo0.0 (Index 64)

```



```
Preferred source address: 203.0.113.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
```

```
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-
  Ethernet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
  PPPoE:
    Dynamic Profile: aci-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    Direct Connect: Off,
    AC Name: nbc
  Input packets : 9
  Output packets: 8
  Protocol multiservice, MTU: Unlimited
```

show interfaces (ALI Interface Set)

```
user@host> show interfaces ge-1/0/0.10
```

```
Logical interface ge-1/0/0.10 (Index 346) (SNMP ifIndex 554) (Generation 155)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.10 ] Encapsulation: ENET2
Line Identity:
  Dynamic Profile: ali-set-profile
  Circuit-id Remote-id Accept-no-ids
  PPPoE:
    Dynamic Profile: ali-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    Direct Connect: Off,
    AC Name: nbc
  Input packets : 9
  Output packets: 8
  Protocol multiservice, MTU: Unlimited
```

Sample Output Gigabit Ethernet

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

user@host> **show interfaces xe-5/0/0 extensive**

```
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 630, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Enabled,
  Flow control: Enabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags       : None
  CoS queues       : 8 supported, 4 maximum usable queues
  Schedulers       : 1024
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53:f6
  Last flapped    : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes   :           6970332384           0 bps
    Output bytes   :                0           0 bps
    Input  packets :           81050506           0 pps
    Output packets :                0           0 pps
  IPv6 transit statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets :                0
    Output packets :                0
  Ingress traffic statistics at Packet Forwarding Engine:
    Input  bytes   :           6970299398           0 bps
    Input  packets :           81049992           0 pps
    Drop  bytes   :                0           0 bps
    Drop  packets :                0           0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
    MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 4 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped packets
```

```

0 best-effort          81049992          81049992          0
1 expedited-fo         0          0          0
2 assured-forw         0          0          0
3 network-cont         0          0          0
Egress queues: 4 supported, 4 in use
Queue counters:        Queued packets  Transmitted packets  Dropped packets
0 best-effort          0          0          0
1 expedited-fo         0          0          0
2 assured-forw         0          0          0
3 network-cont         0          0          0
Active alarms : None
Active defects : None
PCS statistics          Seconds
Bit errors              0
Errored blocks          0
MAC statistics:         Receive      Transmit
Total octets            6970332384      0
Total packets           81050506      0
Unicast packets         81050000      0
Broadcast packets       506          0
Multicast packets       0          0
CRC/Align errors        0          0
FIFO errors              0          0
MAC control frames      0          0
MAC pause frames        0          0
Oversized frames        0
Jabber frames           0
Fragment frames         0
VLAN tagged frames      0
Code violations          0
Filter statistics:
Input packet count      81050506
Input packet rejects    506
Input DA rejects        0
Input SA rejects        0
Output packet count     0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 5
CoS information:
Direction : Output
CoS transmit queue      Bandwidth      Buffer Priority  Limit

```

	%	bps	%	usec		
0 best-effort	95	950000000	95	0	low	none
3 network-control	5	50000000	5	0	low	none
Direction : Input						
CoS transmit queue		Bandwidth		Buffer	Priority	Limit
	%	bps	%	usec		
0 best-effort	95	950000000	95	0	low	none
3 network-control	5	50000000	5	0	low	none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.100] Encapsulation: ENET2

Egress account overhead: 100

Ingress account overhead: 90

Traffic statistics:

Input bytes : 0

Output bytes : 46

Input packets: 0

Output packets: 1

IPv6 transit statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 0

Output bytes : 46

Input packets: 0

Output packets: 1

Transit statistics:

Input bytes : 0 0 bps

Output bytes : 0 0 bps

Input packets: 0 0 pps

Output packets: 0 0 pps

IPv6 transit statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Protocol inet, MTU: 1500, Generation: 253, Route table: 0

Addresses, Flags: Is-Preferred Is-Primary

Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,

Generation: 265

Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0

```
Flags: None
Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

user@host> show interfaces xe-1/0/0 extensive

```
Physical interface: xe-1/0/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 630, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 9.294GbpsGbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled Speed Configuration: Auto
Device flags      : Present Running
Interface flags: SNMP-Traps 16384
Link flags        : None
CoS queues        : 4 supported
Hold-times        : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
Last flapped      : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
  Input  bytes   :                0                0 bps
  Output bytes   :                0                0 bps
  Input  packets :                0                0 pps
  Output packets :                0                0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS Link CRC errors: 0, HS Link FIFO overflows: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
  Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
  Resource errors: 0
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort      0                0                0
  1 expedited-fo     0                0                0
  2 assured-forw     0                0                0
  3 network-cont     0                0                0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics      Seconds      Count
  Bit errors        0                0
  Errored blocks    0                0
```

MAC statistics:	Receive	Transmit
Total octets	0	0
Total packets	0	0
Unicast packets	0	0
Broadcast packets	0	0
Multicast packets	0	0
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	

Filter statistics:

Input packet count	0	
Input packet rejects	0	
Input DA rejects	0	
Input SA rejects	0	
Output packet count		0
Output packet pad count		0
Output packet error count		0

CAM destination filters: 0, CAM source filters: 0

PMA PHY:	Seconds	Count	State
PLL lock	0	0	OK
PHY light	63159	1	Light Missing

WIS section:

BIP-B1	0	0	
SEF	434430	434438	Defect Active
LOS	434430	1	Defect Active
LOF	434430	1	Defect Active
ES-S	434430		
SES-S	434430		
SEFS-S	434430		

WIS line:

BIP-B2	0	0	
REI-L	0	0	
RDI-L	0	0	OK
AIS-L	434430	1	Defect Active
BERR-SF	0	0	OK
BERR-SD	0	0	OK
ES-L	434430		
SES-L	434430		

```

UAS-L          434420
ES-LFE         0
SES-LFE         0
UAS-LFE         0
WIS path:
BIP-B3         0          0
REI-P          0          0
LOP-P          0          0 OK
AIS-P          434430      1 Defect Active
RDI-P          0          0 OK
UNEQ-P         0          0 OK
PLM-P          0          0 OK
ES-P           434430
SES-P           434430
UAS-P           434420
ES-PFE         0
SES-PFE         0
UAS-PFE         0
Received path trace:
 00 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:
  Receive  Transmit
  Total octets 0 0
  Total packets 0 0
  Unicast packets 0 0
  Broadcast packets 0 0
  Multicast packets 0 0
  CRC/Align errors 0 0
  FIFO errors 0 0
  MAC control frames 0 0
  MAC pause frames 0 0

```



```

Oversized frames                0
Jabber frames                   0
Fragment frames                 0
VLAN tagged frames              0
Code violations                  0
Total octets                    0          0
Total packets                   0          0
Unicast packets                 0          0
Broadcast packets               0          0
Multicast packets               0          0
CRC/Align errors                0          0
FIFO errors                     0          0
MAC control frames              0          0
MAC pause frames                0          0
Oversized frames                0
Jabber frames                   0
Fragment frames                 0
VLAN tagged frames              0
Code violations                  0

OTN alarms                      :   None
OTN defects                     :   None
OTN FEC Mode                    : GFEC
OTN Rate                       : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
    Corrected Errors                                0
    Corrected Error Ratio (          0 sec average)    0e-0
OTN FEC alarms:          Seconds          Count  State
    FEC Degrade              0              0  OK
    FEC Excessive            0              0  OK
OTN OC:                  Seconds          Count  State
    LOS                      2              1  OK
    LOF                     67164          2  Defect Active
    LOM                     67164          71  Defect Active
    Wavelength Lock          0              0  OK
OTN OTU:
    AIS                      0              0  OK
    BDI                     65919          4814  Defect Active
    IAE                     67158          1  Defect Active
    TTIM                     7              1  OK
    SF                      67164          2  Defect Active
    SD                      67164          3  Defect Active
    TCA-ES                   0              0  OK
    TCA-SES                   0              0  OK

```

TCA-UAS	80	40	OK
TCA-BBE	0	0	OK
BIP	0	0	OK
BBE	0	0	OK
ES	0	0	OK
SES	0	0	OK
UAS	587	0	OK
Received DAPI:			
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00		
Received SAPI:			
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00		
Transmitted DAPI:			
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00		
Transmitted SAPI:			
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00		
OTN Received Overhead Bytes:			
APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48			
Payload Type: 0x03			
OTN Transmitted Overhead Bytes:			
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00			
Payload Type: 0x03			
Filter statistics:			
Input packet count	0		
Input packet rejects	0		
Input DA rejects	0		
Input SA rejects	0		
Output packet count	0		
Output packet pad count	0		
Output packet error count	0		
CAM destination filters: 0, CAM source filters: 0			
Packet Forwarding Engine configuration:			
Destination slot: 7			
CoS information:			
Direction : Output			
CoS transmit queue	Bandwidth	Buffer Priority	
Limit	%	bps	%
0 best-effort	95	9500000000	95
0			0
low			
none			
3 network-control	5	500000000	5
0			0
low			
none			
...			

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```
user@host> show interfaces xe-7/0/0 extensive
```

```
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional:
  Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags      : Present Running
  ...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```
user@host> show interfaces xe-7/0/0-tx extensive
```

```
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional:
  Tx-Only
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags        : None
  CoS queues        : 8 supported, 8 maximum usable queues
  Hold-times        : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped      : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes   :                0                0 bps
    Output bytes   :      322891152287160      9627472888 bps
    Input  packets :                0                0 pps
    Output packets :      328809727380      1225492 pps
  ...

  Filter statistics:
    Output packet count      328810554250
    Output packet pad count   0
    Output packet error count 0
  ...

  Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)
    Flags: SNMP-Traps Encapsulation: ENET2
    Egress account overhead: 100
```

```

Ingress account overhead: 90
Traffic statistics:
  Input  bytes   :                0
  Output bytes   :      322891152287160
  Input  packets :                0
  Output packets :      328809727380
IPv6 transit statistics:
  Input  bytes   :                0
  Output bytes   :                0
  Input  packets :                0
  Output packets :                0
Local statistics:
  Input  bytes   :                0
  Output bytes   :                0
  Input  packets :                0
  Output packets :                0
Transit statistics:
  Input  bytes   :                0                0 bps
  Output bytes   :      322891152287160          9627472888 bps
  Input  packets :                0                0 pps
  Output packets :      328809727380          1225492 pps
IPv6 transit statistics:
  Input  bytes   :                0
  Output bytes   :                0
  Input  packets :                0
  Output packets :                0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

user@host> show interfaces xe-7/0/0-rx extensive

```

Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional:
Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

```

```

Link flags      : None
CoS queues      : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
  Input  bytes :      322857456303482          9627496104 bps
  Output bytes :                   0              0 bps
  Input  packets:      328775413751          1225495 pps
  Output packets:                   0              0 pps

...

Filter statistics:
  Input packet count      328775015056
  Input packet rejects    1
  Input DA rejects        0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input  bytes :      322857456303482
  Output bytes :                   0
  Input  packets:      328775413751
  Output packets:                   0
IPv6 transit statistics:
  Input  bytes :                   0
  Output bytes :                   0
  Input  packets:                   0
  Output packets:                   0
Local statistics:
  Input  bytes :                   0
  Output bytes :                   0
  Input  packets:                   0
  Output packets:                   0
Transit statistics:
  Input  bytes :      322857456303482          9627496104 bps
  Output bytes :                   0              0 bps
  Input  packets:      328775413751          1225495 pps
  Output packets:                   0              0 pps
IPv6 transit statistics:

```

```

        Input  bytes   :                0
        Output bytes   :                0
        Input  packets :                0
        Output packets :                0
    Protocol inet, MTU: 1500, Generation: 145, Route table: 0
        Addresses, Flags: Is-Preferred Is-Primary
            Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 139
    Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
        Flags: None
        Policer: Input: __default_arp_policer__

```

Sample Output

Sample Output SRX Gigabit Ethernet

user@host> **show interfaces ge-0/0/1**

```

Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: public
  Protocol inet, MTU: 1500

```

```

Flags: 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

```

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: 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

```

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: 0
Mapped forwarding classes
0 best-effort

```



```

1             expedited-forwarding
2             assured-forwarding
3             network-control
Active alarms  : LINK
Active defects : LINK
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes :           0
    Output bytes :           0
    Input  packets:           0
    Output packets:           0
  Local statistics:
    Input  bytes :           0
    Output bytes :           0
    Input  packets:           0
    Output packets:           0
  Transit statistics:
    Input  bytes :           0           0 bps
    Output bytes :           0           0 bps
    Input  packets:           0           0 pps
    Output packets:           0           0 pps
  Security: Zone: public
  Flow Statistics :
  Flow Input statistics :
    Self packets :           0
    ICMP packets :           0
    VPN packets :           0
    Multicast packets :           0
    Bytes permitted by policy :           0
    Connections established :           0
  Flow Output statistics:
    Multicast packets :           0
    Bytes permitted by policy :           0
  Flow error statistics (Packets dropped due to):
    Address spoofing:           0
    Authentication failed:           0
    Incoming NAT errors:           0
    Invalid zone received packet:           0
    Multiple user authentications:           0
    Multiple incoming NAT:           0
    No parent for a gate:           0

```

```

    No one interested in self packets: 0
    No minor session: 0
    No more sessions: 0
    No NAT gate: 0
    No route present: 0
    No SA for incoming SPI: 0
    No tunnel found: 0
    No session for a gate: 0
    No zone or NULL zone binding 0
    Policy denied: 0
    Security association not active: 0
    TCP sequence number out of window: 0
    Syn-attack protection: 0
    User authentication errors: 0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
  Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150

```

show interfaces statistics st0.0 detail

user@host> show interfaces statistics st0.0 detail

```

Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
  Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
  Traffic statistics:
    Input  bytes : 528152756774
    Output bytes : 575950643520
    Input  packets: 11481581669
    Output packets: 12520666095
  Local statistics:
    Input  bytes : 0
    Output bytes : 0
    Input  packets: 0
    Output packets: 0
  Transit statistics:
    Input  bytes : 0 121859888 bps
    Output bytes : 0 128104112 bps
    Input  packets: 0 331141 pps
    Output packets: 0 348108 pps
  Security: Zone: untrust
  Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp

```

```

sap vrrp
Flow Statistics :
Flow Input statistics :
    Self packets :                0
    ICMP packets :                0
    VPN packets :                 0
    Multicast packets :           0
    Bytes permitted by policy :    525984295844
    Connections established :      7
Flow Output statistics:
    Multicast packets :           0
    Bytes permitted by policy :    576003290222
Flow error statistics (Packets dropped due to):
    Address spoofing:             0
    Authentication failed:        0
    Incoming NAT errors:          0
    Invalid zone received packet: 0
    Multiple user authentications: 0
    Multiple incoming NAT:        0
    No parent for a gate:         0
    No one interested in self packets: 0
    No minor session:             0
    No more sessions:             0
    No NAT gate:                  0
    No route present:             2000280
    No SA for incoming SPI:       0
    No tunnel found:              0
    No session for a gate:        0
    No zone or NULL zone binding  0
    Policy denied:                0
    Security association not active: 0
    TCP sequence number out of window: 0
    Syn-attack protection:        0
    User authentication errors:    0
Protocol inet, MTU: 9192
Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
NH drop cnt: 0
Generation: 155, Route table: 0
Flags: Sendbcast-pkt-to-re

```

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1.0 extensive
```

```

Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags      : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues      : 8 supported, 8 maximum usable queues
  Hold-times      : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0                      0 bps
    Output bytes  : 0                      0 bps
    Input packets : 0                      0 pps
    Output packets: 0                      0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort    0                      0                      0
    1 expedited-fo   0                      0                      0
    2 assured-forw   0                      0                      0
    3 network-cont   0                      0                      0
  Queue number:
    Mapped forwarding classes
    0 best-effort
    1 expedited-forwarding
    2 assured-forwarding
    3 network-control
  Active alarms : LINK
  Active defects : LINK
  MAC statistics:
    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             500000000         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			
ge-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

          Packets
Destination class  (packet-per-second)  Bytes
                  (bits-per-second)
          gold      0                    0
                  ( 0 ) ( 0 )
          silver    0                    0
                  ( 0 ) ( 0 )

Logical interface so-0/1/3.0

          Packets
Destination class  (packet-per-second)  Bytes
                  (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
```

Interface	Admin	Link	Proto	Local	Remote Instance
at-0/0/1	up	up	inet	10.0.0.1/24	
ge-0/0/0.0	up	up	inet	192.168.4.28/24	sample-a
at-0/1/0.0	up	up	inet6	fe80::a:0:0:4/64	sample-b
so-0/0/0.0	up	up	inet	10.0.0.1/32	

show interfaces snmp-index (SRX devices)

```
user@host> show interfaces snmp-index 33
```

```

Physical interface: so-2/1/1, Enabled, Physical link is Down
  Interface index: 149, SNMP ifIndex: 33
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : LOL, PLL, LOS
  SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P

```

show interfaces source-class all (SRX devices)

```
user@host> show interfaces source-class all
```

```

Logical interface so-0/1/0.0

```

Source class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	1928095	161959980
(889)	(597762)

```

                bronze                0                0
                (                0) (                0)
                silver                0                0
                (                0) (                0)
Logical interface so-0/1/3.0
                Packets                Bytes
Source class    (packet-per-second)    (bits-per-second)
                gold                0                0
                (                0) (                0)
                bronze                0                0
                (                0) (                0)
                silver                116113            9753492
                (                939) (                631616)

```

show interfaces statistics (Fast Ethernet SRX devices)

user@host> show interfaces fe-1/3/1 statistics

```

Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500
    Flags: Is-Primary, DCU, SCU-in
                Packets                Bytes
Destination class (packet-per-second) (bits-per-second)
                silver1                0                0
                (                0) (                0)
                silver2                0                0
                (                0) (                0)

```



```

                silver3                0                0
                (                0) (                0)
Addresses, Flags: Is-Default Is-Preferred Is-Primary
Destination: 10.27.245/24, Local: 10.27.245.2,
Broadcast: 10.27.245.255
Protocol iso, MTU: 1497
Flags: Is-Primary

```

show interfaces switch-port (SRX devices)

user@host# **show interfaces ge-slot/0/0 switch-port port-number**

```

Port 0, Physical link is Up
Speed: 100mbps, Auto-negotiation: Enabled
Statistics:
    Receive          Transmit
Total bytes          28437086        21792250
Total packets         409145         88008
Unicast packets        9987         83817
Multicast packets     145002           0
Broadcast packets     254156         4191
Multiple collisions    23           10
FIFO/CRC/Align errors 0             0
MAC pause frames       0             0
Oversized frames       0             0
Runt frames            0             0
Jabber frames          0             0
Fragment frames        0             0
Discarded frames       0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
partner Speed: 100 Mbps
Local resolution:
Flow control: None, Remote fault: Link OK

```

show interfaces transport pm (SRX devices)

user@host> **show interfaces transport pm all current et-0/1/0**

```

Physical interface: et-0/1/0, SNMP ifIndex 515
14:45-current          Elapse time:900 Seconds
Near End              Suspect Flag:False          Reason:None

```

PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED	
OTU-BBE	0	800	No	No	
OTU-ES	0	135	No	No	
OTU-SES	0	90	No	No	
OTU-UAS	427	90	No	No	
Far End	Suspect Flag:True		Reason:Unknown		
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED	
OTU-BBE	0	800	No	No	
OTU-ES	0	135	No	No	
OTU-SES	0	90	No	No	
OTU-UAS	0	90	No	No	
Near End	Suspect Flag:False		Reason:None		
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED	
ODU-BBE	0	800	No	No	
ODU-ES	0	135	No	No	
ODU-SES	0	90	No	No	
ODU-UAS	427	90	No	No	
Far End	Suspect Flag:True		Reason:Unknown		
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED	
ODU-BBE	0	800	No	No	
ODU-ES	0	135	No	No	
ODU-SES	0	90	No	No	
ODU-UAS	0	90	No	No	
FEC	Suspect Flag:False		Reason:None		
PM	COUNT	THRESHOLD	TCA-ENABLED	TCA-RAISED	
FEC-CorrectedErr	2008544300	0	NA	NA	
FEC-UncorrectedWords	0	0	NA	NA	
BER	Suspect Flag:False		Reason:None		
PM	MIN	MAX	AVG	THRESHOLD	TCA-ENABLED
TCA-RAISED					
BER	3.6e-5	5.8e-5	3.6e-5	10.0e-3	No
Yes					
Physical interface: et-0/1/0, SNMP ifIndex 515					
14:45-current					
Suspect Flag:True		Reason:Object Disabled			
PM	CURRENT	MIN	MAX	AVG	THRESHOLD
TCA-ENABLED	TCA-RAISED				
					(MIN)
(MAX)	(MIN)	(MAX)	(MIN)	(MAX)	

	Lane chromatic dispersion	0	0	0	0	0
0	NA NA NA NA					
	Lane differential group delay	0	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 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 (Serial)

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 serial interfaces, including RS-232, RS-422/449, EIA-530, X.21, and V.35.

Options

interface-type—On M Series and T Series routers, the interface type is ***se-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 \(Serial, EIA-530\) on page 949](#)

[show interfaces brief \(Serial, EIA-530\) on page 950](#)

[show interfaces detail \(Serial, EIA-530\) on page 950](#)

[show interfaces extensive \(Serial, EIA-530\) on page 951](#)

[show interfaces \(Serial, V.35\) on page 953](#)

[show interfaces brief \(Serial, V.35\) on page 954](#)

[show interfaces detail \(Serial, V.35\) on page 954](#)

[show interfaces extensive \(Serial, V.35\) on page 955](#)

[show interfaces statistics detail \(RS 449\) on page 957](#)

Output Fields

Table 81 on page 942 lists the output fields for the **show interfaces** (Serial) command. Output fields are listed in the approximate order in which they appear.

Table 81: show interfaces (Serial) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 779 .	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU) size on the physical interface.	All levels
Maximum speed	Maximum speed. The nonconfigurable value is 16,384 kbps.	detail extensive none
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 779 .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under “Common Output Fields Description” on page 779 .	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive

Table 81: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Keepalive settings	<p>(PPP and HDLC) Configured settings for keepalive packets.</p> <ul style="list-style-type: none"> • Interval <i>seconds</i>—Time between successive keepalive requests. The range of values, in seconds, is 10 to 32,767. The default value is 10. • Up-count <i>number</i>—Number of keepalive packets a destination must receive to change a link's status from down to up. The range of values is 1 to 255. The default value is 1. • Down-count <i>number</i>—Number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 to 255. The default value is 3. 	All levels
Keepalive	<p>(PPP and HDLC) Information about keepalive packets.</p> <ul style="list-style-type: none"> • Input: <i>number (hh:mm:ss ago)</i>—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • Output: <i>number (hh:mm:ss ago)</i>—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent. 	brief none
Keepalive statistics	<p>(PPP and HDLC) Information about keepalive packets.</p> <ul style="list-style-type: none"> • Input: <i>number (last seen hh:mm:ss ago)</i>—Number of keepalive packets received by PPP and the time since the last keepalive packet was received. • Output: <i>number (last seen hh:mm:ss ago)</i>—Number of keepalive packets sent by PPP and the time since the last keepalive packet was sent. 	detail extensive
LCP state	<p>(PPP) Link Control Protocol state.</p> <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—LCP negotiation is incomplete (not yet completed or has failed). • Not-configured—LCP is not configured on the interface. • Opened—LCP negotiation is successful. 	detail extensive none

Table 81: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
NCP state	<p>(PPP) Network Control Protocol state.</p> <ul style="list-style-type: none"> • Conf-ack-received—Acknowledgement was received. • Conf-ack-sent—Acknowledgement was sent. • Conf-req-sent—Request was sent. • Down—NCP negotiation is incomplete (not yet completed or has failed). • Not-configured—NCP is not configured on the interface. • Opened—NCP negotiation is successful. 	detail extensive none
CHAP state	<p>(PPP) Displays the state of the Challenge Handshake Authentication Protocol (CHAP) during its transaction.</p> <ul style="list-style-type: none"> • Chap-Chal-received—Challenge was received but response not yet sent. • Chap-Chal-sent—Challenge was sent. • Chap-Resp-received—Response was received for the challenge sent, but CHAP has not yet moved into the Success state. (Most likely with RADIUS authentication.) • Chap-Resp-sent—Response was sent for the challenge received. • Closed—CHAP authentication is incomplete. • Failure—CHAP authentication failed. • Not-configured—CHAP is not configured on the interface. • Success—CHAP authentication was successful. 	detail extensive none
CoS queues	Number of CoS queues configured.	detail extensive none
Last flapped	<p>Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago). For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).</p>	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 81: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	extensive

Table 81: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and 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 is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeds the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues supported	Total number of egress queues supported on the specified interface. Displayed with the statistics option.	detail extensive
Egress queues in use	Total number of egress queues in use on the specified interface. Displayed with the statistics option.	detail extensive
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name. Displayed with the statistics option.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 81: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
Serial media information	<p>Information about the physical media:</p> <ul style="list-style-type: none"> • Line protocol—eia530, eia530a, rs232, rs449, v.35, or x.21. • Resync history—Information about resynchronization events: <ul style="list-style-type: none"> • Sync loss count—Number of times the synchronization was lost. • Data signal—(X.21 and V.35) Information about the data signal: <ul style="list-style-type: none"> • Rx Clock—Receive clock status: OK (DTE is receiving the receive clock signal) or Not detected (receive clock signal is not being received). • Control signals—Information about modem control signals: <ul style="list-style-type: none"> • Local mode: DCE (data communication equipment) or DTE (data terminal equipment) • To DCE—Control signals that the Serial PIC sent to the DCE: DTR (Data Terminal Ready: up or down) or RTS (Request To Send: up or down.) • From DC—Control signals that the Serial PIC received from the DCE: CTS (Clear To Send: up or down), DCD (Data Carrier Detect: up or down), DSR (Data Set Ready: up or down), or TM (Test Mode: up or down). • Clocking mode—Clocking used for the transmit clock: <ul style="list-style-type: none"> • dte—Transmit clock is generated by DTE. • dce—Transmit clock is generated by the DCE and is looped back as the transmit clock. • loop-timed—Receive clock from the DCE is looped back as the transmit clock. • Clock rate—Rate, in megahertz (MHz), at which the clock is configured. • Loopback—Configured loopback mode for the interface: dce-remote, dce-local, liu, local, or none. • Tx clock—Clocking phase of the transmit clock: invert (transmit clock polarity is inverted) or non-invert (transmit clock polarity is not inverted). • Line encoding—Type of line encoding used: nrz (nonreturn to zero) or nrzi (return to zero inverted). 	detail extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. • PLP byte—Packet Level Protocol byte. 	extensive

Table 81: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	<p>Information about the CoS queue for the physical interface:</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the source and destination address are also displayed.	brief
Protocol	Protocol family configured on the logical interface, such as iso, inet6, mpls.	detail extensive none

Table 81: show interfaces (Serial) Output Fields (*continued*)

Field Name	Field Description	Level of Output
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 779 .	detail extensive
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (Serial, EIA-530)

user@host> **show interfaces se-5/0/1**

```
Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags   : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags     : Keepalives
```

```

Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 32 (00:00:10 ago), Output: 31 (00:00:07 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
CoS queues      : 8 supported, 8 maximum usable queues
Last flapped    : 2006-04-26 15:10:18 PDT (00:05:22 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3

```

show interfaces brief (Serial, EIA-530)

user@host> **show interfaces se-5/0/1 brief**

```

Physical interface: se-5/0/1, Enabled, Physical link is Up
  Type: Serial, Link-level type: PPP, MTU: 1504
  Device flags      : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags        : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 235 (00:00:10 ago), Output: 234 (00:00:00 ago)

Logical interface se-5/0/1.0
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  inet 12.0.0.1/30

```

show interfaces detail (Serial, EIA-530)

user@host> **show interfaces se-5/0/1 detail**

```

Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41, Generation: 25
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags      : Present Running
  Interface flags: Point-To-Point Internal: 0x4000

```

```

Link flags      : Keepalives
Hold-times      : Up 0 ms, Down 0 ms
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 37 (last seen 00:00:06 ago)
  Output: 35 (last sent 00:00:01 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
CoS queues      : 8 supported, 8 maximum usable queues
Last flapped    : 2006-04-26 15:10:18 PDT (00:06:02 ago)
Statistics last cleared: Never
Traffic statistics:
  Input  bytes :                928                40 bps
  Output bytes :               1023                48 bps
  Input  packets:                 76                 0 pps
  Output packets:                 77                 0 pps
Serial media information:
  Line protocol: eia530
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: OK
  Control signals:
    Local mode: DTE
    To DCE: DTR: up, RTS: up
    From DCE: CTS: up, DCD: up, DSR: up
  Clocking mode: loop-timed
  Clock rate: 8.0 MHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Protocol inet, MTU: 1500, Generation: 15, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
    Generation: 23

```

show interfaces extensive (Serial, EIA-530)

```
user@host> show interfaces se-5/0/1 extensive
```

```

Physical interface: se-5/0/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 41, Generation: 25
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags      : Present Running
  Interface flags: Point-To-Point Internal: 0x4000
  Link flags        : Keepalives
  Hold-times        : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 40 (last seen 00:00:00 ago)
    Output: 37 (last sent 00:00:09 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues      : 8 supported, 8 maximum usable queues
  Last flapped    : 2006-04-26 15:10:18 PDT (00:06:28 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes   :                988                40 bps
    Output bytes   :               1088                48 bps
    Input  packets :                81                 0 pps
    Output packets :                82                 0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 2, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
  Serial media information:
    Line protocol: eia530
    Resync history:
      Sync loss count: 0
    Data signal:
      Rx Clock: OK
    Control signals:
      Local mode: DTE
      To DCE: DTR: up, RTS: up
      From DCE: CTS: up, DCD: up, DSR: up
    Clocking mode: loop-timed
    Clock rate: 8.0 MHz
    Loopback: none
    Tx clock: non-invert
    Line encoding: nrz

```



```

Packet Forwarding Engine configuration:
  Destination slot: 5, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue      Bandwidth      Buffer  Priority  Limit
                           %      bps      %      usec
  0 best-effort           95      15564800  95        0      low   none
  3 network-control       5       819200   5         0      low   none

Logical interface se-5/0/1.0 (Index 71) (SNMP ifIndex 45) (Generation 9)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Protocol inet, MTU: 1500, Generation: 15, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 12.0.0.0/30, Local: 12.0.0.1, Broadcast: 12.0.0.3,
    Generation: 23

```

show interfaces (Serial, V.35)

user@host> **show interfaces se-5/0/0**

```

Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags      : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags       : Loose-NCP
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 0 (never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
  CoS queues      : 8 supported, 8 maximum usable queues
  Last flapped    : 2006-04-26 14:51:27 PDT (01:02:23 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27)
  Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
  Encapsulation: PPP
  Protocol inet, MTU: 1500
  Flags: Protocol-Down
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3

```

show interfaces brief (Serial, V.35)

```
user@host> show interfaces se-5/0/0 brief
```

```
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Type: Serial, Link-level type: PPP, MTU: 1504
  Device flags      : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags        : Loose-NCP
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 0 (never)

Logical interface se-5/0/0.0
  Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
  Encapsulation: PPP
  inet 13.0.0.2/30
```

show interfaces detail (Serial, V.35)

```
user@host> show interfaces se-5/0/0 detail
```

```
Physical interface: se-5/0/0, Enabled, Physical link is Down
  Interface index: 150, SNMP ifIndex: 39, Generation: 31
  Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
  Device flags      : Present Running Down
  Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
  Link flags        : Loose-NCP
  Hold-times        : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive statistics:
    Input : 0 (last seen: never)
    Output: 0 (last sent: never)
  LCP state: Down
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Closed
  CoS queues        : 8 supported, 8 maximum usable queues
  Last flapped      : 2006-04-26 14:51:27 PDT (01:03:15 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
  Serial media information:
```

```

Line protocol: v.35
Resync history:
  Sync loss count: 0
Data signal:
  Rx Clock: Not Detected
Control signals:
  Local mode: DCE
  To DTE: CTS: down, DCD: down, DSR: up
  From DTE: DTR: down, RTS: down
DCE loopback override: Off
Clocking mode: internal
Clock rate: 38.4 KHz
Loopback: none
Tx clock: non-invert
Line encoding: nrz

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Protocol inet, MTU: 1500, Generation: 17, Route table: 0
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,
  Generation: 23

```

show interfaces extensive (Serial, V.35)

user@host> **show interfaces se-5/0/0 extensive**

```

Physical interface: se-5/0/0, Enabled, Physical link is Down
Interface index: 150, SNMP ifIndex: 39, Generation: 31
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 16384kbps
Device flags      : Present Running Down
Interface flags: Hardware-Down Point-To-Point Internal: 0x4000
Link flags       : Loose-NCP
Hold-times       : Up 0 ms, Down 0 ms
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 0 (last seen: never)
  Output: 0 (last sent: never)
LCP state: Down
NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Closed

```

```

CoS queues      : 8 supported, 8 maximum usable queues
Last flapped   : 2006-04-26 14:51:27 PDT (01:04:17 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :                0                0 bps
  Output bytes  :                0                0 bps
  Input packets :                0                0 pps
  Output packets:                0                0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
Serial media information:
  Line protocol: v.35
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: Not Detected
  Control signals:
    Local mode: DCE
    To DTE: CTS: down, DCD: down, DSR: up
    From DTE: DTR: down, RTS: down
  DCE loopback override: Off
  Clocking mode: internal
  Clock rate: 38.4 KHz
  Loopback: none
  Tx clock: non-invert
  Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 5, PLP byte: 1 (0x00)
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %          bps      %      usec
  0 best-effort           95        15564800  95        0        low    none
  3 network-control        5         819200   5         0        low    none

Logical interface se-5/0/0.0 (Index 73) (SNMP ifIndex 27) (Generation 12)
Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Protocol inet, MTU: 1500, Generation: 17, Route table: 0
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

```

```
Destination: 13.0.0.0/30, Local: 13.0.0.2, Broadcast: 13.0.0.3,
Generation: 23
```

show interfaces statistics detail (RS 449)

```
user@host> show interfaces se-6/0/0 statistics detail
```

```
Interface index: 149, SNMP ifIndex: 59, Generation: 150
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 8mbps
Device flags      : Present Running
Interface flags: Point-To-Point Internal: 0x4000
Link flags        : No-Keepalives Loose-NCP
Hold-times        : Up 0 ms, Down 0 ms
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
CoS queues        : 8 supported, 8 maximum usable queues
Last flapped      : 2007-11-28 19:38:36 PST (00:14:06 ago)
Statistics last cleared: Never
Traffic statistics:
  Input  bytes   :                744                0 bps
  Output bytes   :               5978                0 bps
  Input  packets :                33                0 pps
  Output packets :               129                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: 13, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Egress queues: 8 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort      24                24                0
  1 expedited-fo      0                0                0
  2 bulk              0                0                0
  3 assured-forw     105               105                0
  4 voip              0                0                0
Serial media information:
  Line protocol: rs449
  Resync history:
    Sync loss count: 0
```

```
Data signal:
  Rx Clock: OK
Control signals:
  Local mode: DTE
  To DCE: DTR: up, RTS: up
  From DCE: CTS: up, DCD: up, DSR: up
Clocking mode: internal
Loopback: none
Tx clock: non-invert
Line encoding: nrz
```

```
Logical interface se-6/0/0.0 (Index 75) (SNMP ifIndex 69) (Generation 141)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
Protocol inet, MTU: 256, Generation: 145, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 11.11.11/24, Local: 11.11.11.2, Broadcast: 11.11.11.255,
  Generation: 157
```

show interfaces diagnostics optics

Syntax

```
show interfaces diagnostics optics interface-name
```

Release Information

Command introduced in Junos OS Release 10.0 for EX Series switches.

Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, XFP, QSFP+, or CFP) installed in EX Series or QFX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

Options

interface-name—Name of the interface associated with the port in which the transceiver is installed:

ge-fpc/pic/port, *xe-fpc/pic/port*, or *et-fpc/pic/port*.

Required Privilege Level

view

RELATED DOCUMENTATION

Monitoring Interface Status and Traffic

[Monitoring Interface Status and Traffic | 392](#)

Install a Transceiver

Installing a Transceiver in a QFX Series Device

Remove a Transceiver

Removing a Transceiver from a QFX Series Device

[Junos OS Ethernet Interfaces Configuration Guide](#)

List of Sample Output

[show interfaces diagnostics optics ge-0/1/0 \(SFP Transceiver\) on page 980](#)

[show interfaces diagnostics optics xe-0/1/0 \(SFP+ Transceiver\) on page 981](#)

[show interfaces diagnostics optics xe-0/1/0 \(XFP Transceiver\) on page 983](#)

[show interfaces diagnostics optics et-3/0/0 \(QSFP+ Transceiver\) on page 984](#)

[show interfaces diagnostics optics et-4/1/0 \(CFP Transceiver\) on page 985](#)

Output Fields

Table 82 on page 960 lists the output fields for the **show interfaces diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

Transceivers generate monitoring data such as power supply voltage, transceiver temperature, TEC current, receive optical power, laser output power, laser temperature, and laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.

Warning flags associated with power supply voltage, transceiver temperature, TEC current, receive optical power, laser output power, laser temperature, and laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.

Alarm flags associated with power supply voltage, transceiver temperature, TEC current, receive optical power, laser output power, laser temperature, and laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.

Table 82: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Lane 'x'	Displays an individual Tx/Rx data transmission channel "x" associated with the given physical interface device (or IFD).
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power (Not available for QSFP+ transceivers)	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser temperature (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the laser temperature, in Celsius and Fahrenheit.
Module temperature	Displays the temperature, in Celsius and Fahrenheit.

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage (Not available for XFP transceivers)	Displays the voltage, in Volts.
Laser rx power (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Receiver signal average optical power (Not available for XFP, QSFP+, and CFP transceivers)	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	<p>Displays whether the laser bias power setting high alarm is On or Off.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser bias current low alarm	<p>Displays whether the laser bias power setting low alarm is On or Off.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser bias current high warning	<p>Displays whether the laser bias power setting high warning is On or Off.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser bias current low warning	<p>Displays whether the laser bias power setting low warning is On or Off.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser output power high alarm (Not available for QSFP+ transceivers)	<p>Displays whether the laser output power high alarm is On or Off.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser output power low alarm (Not available for QSFP+ transceivers)	<p>Displays whether the laser output power low alarm is On or Off.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser output power high warning (Not available for QSFP+ transceivers)	<p>Displays whether the laser output power high warning is On or Off.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser output power low warning (Not available for QSFP+ transceivers)	<p>Displays whether the laser output power low warning is On or Off.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser temperature high alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays whether the laser temperature high alarm is On or Off.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser temperature low alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays whether the laser temperature low alarm is On or Off.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser temperature high warning (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays whether the laser temperature high warning is On or Off.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser temperature low warning (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays whether the laser temperature low warning is On or Off.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Module temperature high alarm (Not available for QSFP+ transceivers)	<p>Displays whether the module temperature high alarm is On or Off.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Module temperature low alarm (Not available for QSFP+ transceivers)	<p>Displays whether the module temperature low alarm is On or Off.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module temperature high warning (Not available for QSFP+ transceivers)	<p>Displays whether the module temperature high warning is On or Off.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Module temperature low warning (Not available for QSFP+ transceivers)	<p>Displays whether the module temperature low warning is On or Off.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Module voltage high alarm (Not available for XFP and QSFP+ transceivers)	<p>Displays whether the module voltage high alarm is On or Off.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage low alarm (Not available for XFP and QSFP+ transceivers)	<p>Displays whether the module voltage low alarm is On or Off.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Module voltage high warning (Not available for XFP and QSFP+ transceivers)	<p>Displays whether the module voltage high warning is On or Off.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Module voltage low warning (Not available for XFP and QSFP+ transceivers)	<p>Displays whether the module voltage low warning is On or Off.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power high alarm (Not available for QSFP+ and CFP transceivers)	<p>Displays whether the receive laser power high alarm is On or Off.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser rx power low alarm (Not available for QSFP+ and CFP transceivers)	<p>Displays whether the receive laser power low alarm is On or Off.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser rx power high warning (Not available for QSFP+ and CFP transceivers)	<p>Displays whether the receive laser power high warning is On or Off.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power low warning (Not available for QSFP+ and CFP transceivers)	<p>Displays whether the receive laser power low warning is On or Off.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser bias current high alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser bias current high alarm.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Module not ready alarm (Not available for SFP, SFP+, and QSFP+ transceivers)	<p>Displays whether the module not ready alarm is On or Off. When the output is On, the module has an operational fault.</p>
Module low power alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays whether the module low power alarm is On or Off.</p> <p>The host may use the module power mode to limit module power consumption to one of two thresholds—low power mode or high power mode. When operating in low power mode, the maximum power consumption of the module will remain below a defined maximum. Active module functionality is limited to host-to-module interactions over the management interface. Passive copper cable assemblies may operate in low power mode at all times and are fully functional.</p>
Module initialization incomplete alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays whether the module initialization incomplete alarm is On or Off.</p> <p>This alarm is generated when the transceiver module initialization is not complete.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module fault alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the module fault alarm is On or Off . This alarm is generated when the transceiver module detects a condition (such as TEC runaway or flash corruption) that could cause damage to the module. This alarm is provided as a notification to the host that a module fault has occurred. The implementation of the module fault is form factor dependent.
PLD Flash initialization fault alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the PLD Flash initialization fault alarm is On or Off . This alarm is generated if the initialization of transceiver module internal programmable logic device (PLD) or flash device fails.
Power supply fault alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the power supply fault alarm is On or Off . This alarm is generated if the power supply is out of range.
Checksum fault alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the checksum fault alarm is On or Off . To protect against degradation in the non-volatile storage and potential errors in transfer, the module management interface defines a few registers to let module manufacturers save checksums or check codes. Each checksum or check code is assigned for a small-sized datum derived from a block of digital data. The algorithms of calculation for checksums or check codes are defined by management interface specifications of form factors. For example, the checksum or check code shall be the low order 8 bits of the sum of the contents of that small-sized datum. Prior to using the stored information, module users calculate the checksums or check codes again using the same algorithm. If the results are different from those saved in checksum bytes by module manufacturers, checksum fault alarm is triggered. A Checksum fault means that the content of the small-sized datum covered by this checksum is corrupted or changed.
Tx laser disabled alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the Tx laser disabled alarm is On or Off . This alarm is generated when the transmitter optical output is shut down.
Module power down alarm (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Displays whether the module power down alarm is On or Off . When the output is On , module is in a limited power mode, low for normal operation.

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Tx data not ready alarm (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is On or Off .
Tx not ready alarm (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is On or Off .
Tx laser fault alarm (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Laser fault condition. Displays whether the Tx laser fault alarm is On or Off . This alarm is generated when the transmitter/laser is operating in an abnormal state.
Tx CDR loss of lock alarm (Not available for SFP, SFP+, and QSFP+ transceivers)	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is On or Off . Some transceivers have Clock and Data Recovery (CDR) circuits at the transmit side. The CDR has its own oscillator and uses a phase locked loop (PLL) to calibrate the oscillator and match it to the phase and frequency of the incoming signal. The feedback circuit of PLL eventually locks the local oscillator on to the phase and frequency of the incoming signal. The PLL maintains that lock by continuously monitoring the received signal. If the lock cannot be maintained, the Loss of Lock (LOL) alarm is generated.
Rx not ready alarm (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is On or Off .
Rx loss of signal alarm (Not available for SFP and SFP+ transceivers)	Receive loss of signal alarm. When the output is On , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is On or Off .

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Rx CDR loss of lock alarm (Not available for SFP, SFP+, and QSFP+ transceivers)	<p>Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is On or Off.</p> <p>Some transceivers have Clock and Data Recovery (CDR) circuits at the receive side. The CDR has its own oscillator and uses a phase locked loop (PLL) to calibrate the oscillator and match it to the phase and frequency of the incoming signal. The feedback circuit of PLL eventually locks the local oscillator on to the phase and frequency of the incoming signal. The PLL maintains that lock by continuously monitoring the received signal. If the lock cannot be maintained, the Loss of Lock (LOL) alarm is generated.</p>
Laser bias current low alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser bias current low alarm.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser bias current indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser bias current high warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser bias current high warning.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser bias current low warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser bias current low warning.</p> <p>Transceivers generate monitoring data such as laser bias current by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser bias current indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser output power high alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser output power high alarm.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser output power low alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser output power low alarm.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser output power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser output power high warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser output power high warning.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser output power low warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser output power low warning.</p> <p>Transceivers generate monitoring data such as laser output power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser output power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Module temperature high alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module temperature high alarm.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module temperature low alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module temperature low alarm.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with laser temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Module temperature high warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module temperature high warning.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Module temperature low warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module temperature low warning.</p> <p>Transceivers generate monitoring data such as laser temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with laser temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage high alarm threshold (Not available for XFP and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module voltage high alarm.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Module voltage low alarm threshold (Not available for XFP and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module voltage low alarm.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with power supply voltage indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Module voltage high warning threshold (Not available for XFP and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module voltage high warning.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage low warning threshold (Not available for XFP and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the module voltage low warning.</p> <p>Transceivers generate monitoring data such as power supply voltage by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with power supply voltage indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser rx power high alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser rx power high alarm.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser rx power low alarm threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser rx power low alarm.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with receive optical power indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
Laser rx power high warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser rx power high warning.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser rx power low warning threshold (Not available for QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser rx power low warning.</p> <p>Transceivers generate monitoring data such as receive optical power by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with receive optical power indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser temperature high alarm threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser temperature high alarm, in Celsius and Fahrenheit.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser temperature low alarm threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser temperature low alarm, in Celsius and Fahrenheit.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Alarm flags associated with transceiver temperature indicate conditions likely, but not necessarily, to be associated with an operational failure of the transceiver and cause for immediate action.</p>
Laser temperature high warning threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser temperature high warning, in Celsius and Fahrenheit.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
Laser temperature low warning threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for the laser temperature low warning, in Celsius and Fahrenheit.</p> <p>Transceivers generate monitoring data such as transceiver temperature by digitization of internal analog signals. Each supported analog to digital conversion value has a high alarm, low alarm, high warning, and low warning threshold. These values are configured by the transceiver manufacturer and allow the user to determine when a value is outside the normal limits.</p> <p>Warning flags associated with transceiver temperature indicate conditions outside the normal limits; however, that need not cause immediate link failures.</p>
SOA bias current high alarm threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	<p>Displays the vendor-specified threshold for SOA bias current high alarm.</p>

Table 82: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
SOA bias current low alarm threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for SOA bias current low alarm.
SOA bias current high warning threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for SOA bias current high warning.
SOA bias current low warning threshold (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for SOA bias current low warning.
Laser receiver power high alarm (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power high alarm is On or Off . This alarm is triggered when the receive optical power is higher than the high alarm threshold.
Laser receiver power low alarm (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power low alarm is On or Off . This alarm is triggered when the receive optical power is lower than the high alarm threshold.
Laser receiver power high warning (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power high warning is On or Off . This warning is triggered when the receive optical power is higher than the high warning threshold.
Laser receiver power low warning (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power low warning is On or Off . This warning is triggered when the receive optical power is lower than the high warning threshold.
Laser receiver power (Not available for SFP, SFP+, and XFP transceivers)	Displays the laser receiver power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).

Table 82: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
Tx loss of signal functionality alarm (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the Tx loss of signal functionality alarm is On or Off . This alarm is generated if the high-speed electrical signal amplitude of Tx input is below a threshold.
APD supply fault alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the APD supply fault alarm is On or Off . Some transceivers have Avalanche Photodiode (APD) receiver to improve receiver sensitivity. APD requires a bias voltage. If there is a disruption in the bias voltage supply, the APD supply fault alarm is generated.
TEC fault alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the TEC fault alarm is On or Off . Some transceivers have an internal temperature controller—Thermoelectric Cooler (TEC). If the TEC is faulty, the TEC fault alarm is generated.
Wavelength unlocked alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the Wavelength unlocked alarm is On or Off . The transmitter wavelengths of some transceivers must be maintained at greater accuracy over operational lifetimes. Wavelength locker is used to stabilize the wavelength of the laser diode. If the wavelength is out of the specified range, the wavelength unlocked alarm is generated.

Sample Output

show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver)

```
user@switch> show interfaces diagnostics optics ge-0/1/0
```

```
Physical interface: ge-0/1/0
  Laser bias current           : 5.444 mA
  Laser output power          : 0.3130 mW / -5.04 dBm
  Module temperature          : 36 degrees C / 97 degrees F
  Module voltage               : 3.2120 V
  Receiver signal average optical power : 0.3840 mW / -4.16 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
```

```

Laser output power low alarm           : Off
Laser output power high warning        : Off
Laser output power low warning         : Off
Module temperature high alarm          : Off
Module temperature low alarm           : Off
Module temperature high warning        : Off
Module temperature low warning         : Off
Module voltage high alarm              : Off
Module voltage low alarm               : Off
Module voltage high warning            : Off
Module voltage low warning             : Off
Laser rx power high alarm              : Off
Laser rx power low alarm               : Off
Laser rx power high warning            : Off
Laser rx power low warning             : Off
Laser bias current high alarm threshold : 15.000 mA
Laser bias current low alarm threshold  : 1.000 mA
Laser bias current high warning threshold : 12.000 mA
Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6300 mW / -2.01 dBm
Laser output power low alarm threshold  : 0.0660 mW / -11.80 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0780 mW / -11.08 dBm
Module temperature high alarm threshold : 109 degrees C / 228 degrees F
Module temperature low alarm threshold  : -29 degrees C / -20 degrees F
Module temperature high warning threshold : 103 degrees C / 217 degrees F
Module temperature low warning threshold : -13 degrees C / 9 degrees F
Module voltage high alarm threshold     : 3.900 V
Module voltage low alarm threshold      : 2.700 V
Module voltage high warning threshold   : 3.700 V
Module voltage low warning threshold    : 2.900 V
Laser rx power high alarm threshold     : 1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold      : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold   : 0.7939 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.0157 mW / -18.04 dBm

```

Sample Output

show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver)

```
user@switch> show interfaces diagnostics optics xe-0/1/0
```

Physical interface: xe-0/1/0

Laser bias current	: 4.968 mA
Laser output power	: 0.4940 mW / -3.06 dBm
Module temperature	: 27 degrees C / 81 degrees F
Module voltage	: 3.2310 V
Receiver signal average optical power	: 0.0000
Laser bias current high alarm	: Off
Laser bias current low alarm	: Off
Laser bias current high warning	: Off
Laser bias current low warning	: Off
Laser output power high alarm	: Off
Laser output power low alarm	: Off
Laser output power high warning	: Off
Laser output power low warning	: Off
Module temperature high alarm	: Off
Module temperature low alarm	: Off
Module temperature high warning	: Off
Module temperature low warning	: Off
Module voltage high alarm	: Off
Module voltage low alarm	: Off
Module voltage high warning	: Off
Module voltage low warning	: Off
Laser rx power high alarm	: Off
Laser rx power low alarm	: On
Laser rx power high warning	: Off
Laser rx power low warning	: On
Laser bias current high alarm threshold	: 10.500 mA
Laser bias current low alarm threshold	: 2.000 mA
Laser bias current high warning threshold	: 9.000 mA
Laser bias current low warning threshold	: 2.500 mA
Laser output power high alarm threshold	: 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold	: 0.0740 mW / -11.31 dBm
Laser output power high warning threshold	: 0.7070 mW / -1.51 dBm
Laser output power low warning threshold	: 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold	: 75 degrees C / 167 degrees F
Module temperature low alarm threshold	: -5 degrees C / 23 degrees F
Module temperature high warning threshold	: 70 degrees C / 158 degrees F
Module temperature low warning threshold	: 0 degrees C / 32 degrees F
Module voltage high alarm threshold	: 3.630 V
Module voltage low alarm threshold	: 2.970 V
Module voltage high warning threshold	: 3.465 V
Module voltage low warning threshold	: 3.135 V
Laser rx power high alarm threshold	: 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold	: 0.0407 mW / -13.90 dBm

```

Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold  : 0.1023 mW / -9.90 dBm

```

Sample Output

show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver)

user@switch> **show interfaces diagnostics optics xe-0/1/0**

```

Physical interface: xe-0/1/0
  Laser bias current           : 8.029 mA
  Laser output power           : 0.6430 mW / -1.92 dBm
  Module temperature           : 4 degrees C / 39 degrees F
  Laser rx power               : 0.0012 mW / -29.21 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm  : Off
  Laser output power low alarm   : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm  : Off
  Module temperature low alarm   : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Laser rx power high alarm      : Off
  Laser rx power low alarm       : On
  Laser rx power high warning    : Off
  Laser rx power low warning     : On
  Module not ready alarm         : On
  Module power down alarm        : Off
  Tx data not ready alarm        : Off
  Tx not ready alarm             : Off
  Tx laser fault alarm           : Off
  Tx CDR loss of lock alarm      : Off
  Rx not ready alarm             : On
  Rx loss of signal alarm        : On
  Rx CDR loss of lock alarm      : On
  Laser bias current high alarm threshold : 13.000 mA
  Laser bias current low alarm threshold  : 2.000 mA
  Laser bias current high warning threshold : 12.000 mA

```

```

Laser bias current low warning threshold : 3.000 mA
Laser output power high alarm threshold  : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold   : 0.1650 mW / -7.83 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 90 degrees C / 194 degrees F
Module temperature low alarm threshold    : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold       : 0.8912 mW / -0.50 dBm
Laser rx power low alarm threshold        : 0.0912 mW / -10.40 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold      : 0.1023 mW / -9.90 dBm

```

Sample Output

show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver)

user@switch> **show interfaces diagnostics optics et-3/0/0**

```

Physical interface: et-3/0/0
  Module temperature           : 33 degrees C / 92 degrees F
  Module voltage               : 3.3060 V
  Lane 0
    Laser bias current         : 7.182 mA
    Laser receiver power       : 0.743 mW / -1.29 dBm
    Laser bias current high alarm : Off
    Laser bias current low alarm  : Off
    Laser bias current high warning : Off
    Laser bias current low warning : Off
    Laser receiver power high alarm : Off
    Laser receiver power low alarm  : Off
    Laser receiver power high warning : Off
    Laser receiver power low warning : Off
    Tx loss of signal functionality alarm : Off
    Rx loss of signal alarm        : Off
  Lane 1
    Laser bias current         : 7.326 mA
    Laser receiver power       : 0.752 mW / -1.24 dBm
    Laser bias current high alarm : Off
    Laser bias current low alarm  : Off
    Laser bias current high warning : Off

```



```

Laser bias current low warning           : Off
Laser receiver power high alarm          : Off
Laser receiver power low alarm           : Off
Laser receiver power high warning        : Off
Laser receiver power low warning         : Off
Tx loss of signal functionality alarm     : Off
Rx loss of signal alarm                   : Off
Lane 2
Laser bias current                       : 7.447 mA
Laser receiver power                     : 0.790 mW / -1.03 dBm
Laser bias current high alarm             : Off
Laser bias current low alarm              : Off
Laser bias current high warning           : Off
Laser bias current low warning            : Off
Laser receiver power high alarm           : Off
Laser receiver power low alarm            : Off
Laser receiver power high warning         : Off
Laser receiver power low warning          : Off
Tx loss of signal functionality alarm      : Off
Rx loss of signal alarm                   : Off
Lane 3
Laser bias current                       : 7.734 mA
Laser receiver power                     : 0.768 mW / -1.15 dBm
Laser bias current high alarm             : Off
Laser bias current low alarm              : Off
Laser bias current high warning           : Off
Laser bias current low warning            : Off
Laser receiver power high alarm           : Off
Laser receiver power low alarm            : Off
Laser receiver power high warning         : Off
Laser receiver power low warning          : Off
Tx loss of signal functionality alarm      : Off
Rx loss of signal alarm                   : Off

```

Sample Output

show interfaces diagnostics optics et-4/1/0 (CFP Transceiver)

user@switch> **show interfaces diagnostics optics et-4/1/0**

```

Physical interface: et-4/1/0
Module temperature           : 38 degrees C / 101 degrees F

```

```

Module voltage : 3.2500 V
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Module not ready alarm : Off
Module low power alarm : Off
Module initialization incomplete alarm : Off
Module fault alarm : Off
PLD Flash initialization fault alarm : Off
Power supply fault alarm : Off
Checksum fault alarm : Off
Tx laser disabled alarm : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm : Off
Rx loss of signal alarm : Off
Rx CDR loss of lock alarm : Off
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.5000 V
Module voltage low alarm threshold : 3.0990 V
Module voltage high warning threshold : 3.4000 V
Module voltage low warning threshold : 3.2000 V
Laser bias current high alarm threshold : 250.000 mA
Laser bias current low alarm threshold : 37.500 mA
Laser bias current high warning threshold : 225.000 mA
Laser bias current low warning threshold : 50.000 mA
Laser output power high alarm threshold : 3.9800 mW / 6.00 dBm
Laser output power low alarm threshold : 0.4670 mW / -3.31 dBm
Laser output power high warning threshold : 3.5480 mW / 5.50 dBm
Laser output power low warning threshold : 0.5240 mW / -2.81 dBm
Laser rx power high alarm threshold : 3.5481 mW / 5.50 dBm
Laser rx power low alarm threshold : 0.0616 mW / -12.10 dBm
Laser rx power high warning threshold : 3.1622 mW / 5.00 dBm
Laser rx power low warning threshold : 0.0691 mW / -11.61 dBm
Laser temperature high alarm threshold : 67 degrees C / 153 degrees F
Laser temperature low alarm threshold : 35 degrees C / 95 degrees F
Laser temperature high warning threshold : 62 degrees C / 144 degrees F

```

```

Laser temperature low warning threshold : 40 degrees C / 104 degrees F
SOA bias current high alarm threshold  : 0.000 mA
SOA bias current low alarm threshold    : 0.000 mA
SOA bias current high warning threshold : 0.000 mA
SOA bias current low warning threshold  : 0.000 mA
Lane 0
Laser bias current                      : 131.684 mA
Laser output power                      : 1.002 mW / 0.01 dBm
Laser temperature                      : 54 degrees C / 128 degrees F
Laser receiver power                   : 0.497 mW / -3.03 dBm
Laser bias current high alarm          : Off
Laser bias current low alarm           : Off
Laser bias current high warning        : Off
Laser bias current low warning         : Off
Laser output power high alarm          : Off
Laser output power low alarm           : Off
Laser output power high warning        : Off
Laser output power low warning         : Off
Laser temperature high alarm           : Off
Laser temperature low alarm            : Off
Laser temperature high warning         : Off
Laser temperature low warning          : Off
Laser receiver power high alarm        : Off
Laser receiver power low alarm         : Off
Laser receiver power high warning      : Off
Laser receiver power low warning       : Off
Tx loss of signal functionality alarm  : Off
Tx CDR loss of lock alarm              : Off
Rx loss of signal alarm                : Off
Rx CDR loss of lock alarm              : Off
APD supply fault alarm                 : Off
TEC fault alarm                        : Off
Wavelength unlocked alarm              : Off
Lane 1
Laser bias current                      : 122.345 mA
Laser output power                      : 1.002 mW / 0.01 dBm
Laser temperature                      : 51 degrees C / 124 degrees F
Laser receiver power                   : 0.611 mW / -2.14 dBm
Laser bias current high alarm          : Off
Laser bias current low alarm           : Off
Laser bias current high warning        : Off
Laser bias current low warning         : Off
Laser output power high alarm          : Off
Laser output power low alarm           : Off

```

Laser output power high warning	: Off
Laser output power low warning	: Off
Laser temperature high alarm	: Off
Laser temperature low alarm	: Off
Laser temperature high warning	: Off
Laser temperature low warning	: Off
Laser receiver power high alarm	: Off
Laser receiver power low alarm	: Off
Laser receiver power high warning	: Off
Laser receiver power low warning	: Off
Tx loss of signal functionality alarm	: Off
Tx CDR loss of lock alarm	: Off
Rx loss of signal alarm	: Off
Rx CDR loss of lock alarm	: Off
APD supply fault alarm	: Off
TEC fault alarm	: Off
Wavelength unlocked alarm	: Off
Lane 2	
Laser bias current	: 112.819 mA
Laser output power	: 1.000 mW / 0.00 dBm
Laser temperature	: 50 degrees C / 122 degrees F
Laser receiver power	: 0.540 mW / -2.67 dBm
Laser bias current high alarm	: Off
Laser bias current low alarm	: Off
Laser bias current high warning	: Off
Laser bias current low warning	: Off
Laser output power high alarm	: Off
Laser output power low alarm	: Off
Laser output power high warning	: Off
Laser output power low warning	: Off
Laser temperature high alarm	: Off
Laser temperature low alarm	: Off
Laser temperature high warning	: Off
Laser temperature low warning	: Off
Laser receiver power high alarm	: Off
Laser receiver power low alarm	: Off
Laser receiver power high warning	: Off
Laser receiver power low warning	: Off
Tx loss of signal functionality alarm	: Off
Tx CDR loss of lock alarm	: Off
Rx loss of signal alarm	: Off
Rx CDR loss of lock alarm	: Off
APD supply fault alarm	: Off
TEC fault alarm	: Off

Wavelength unlocked alarm	: Off
Lane 3	
Laser bias current	: 100.735 mA
Laser output power	: 1.002 mW / 0.01 dBm
Laser temperature	: 50 degrees C / 122 degrees F
Laser receiver power	: 0.637 mW / -1.96 dBm
Laser bias current high alarm	: Off
Laser bias current low alarm	: Off
Laser bias current high warning	: Off
Laser bias current low warning	: Off
Laser output power high alarm	: Off
Laser output power low alarm	: Off
Laser output power high warning	: Off
Laser output power low warning	: Off
Laser temperature high alarm	: Off
Laser temperature low alarm	: Off
Laser temperature high warning	: Off
Laser temperature low warning	: Off
Laser receiver power high alarm	: Off
Laser receiver power low alarm	: Off
Laser receiver power high warning	: Off
Laser receiver power low warning	: Off
Tx loss of signal functionality alarm	: Off
Tx CDR loss of lock alarm	: Off
Rx loss of signal alarm	: Off
Rx CDR loss of lock alarm	: Off
APD supply fault alarm	: Off
TEC fault alarm	: Off
Wavelength unlocked alarm	: Off

show interfaces extensive

Syntax

```
show interfaces extensive
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 17.2 for PT1000 and PTX10008 Packet Transport Routers.

Description

Display extensive information about all interfaces configured on the router.

NOTE:

- At some times, the cumulative byte counters displayed with the **show interfaces extensive** command on the 10-Gigabit Ethernet MPC with SFP+ is not always increasing and cumulative and does not give the correct results. There is a time lag in collecting these statistics, during which the display might decrease or go from a nonzero number to zero. Eventually, the counter will display the correct result.
- When the **show interfaces extensive** command is executed on a router with an MPC or a T4000 Type 5 FPC, the **Input packet rejects** counter of the **Filter statistics** field also displays statistics related to the following packet errors:
 - Invalid VLAN range
 - Tagged packet received on an untagged interface
- When the **show interfaces extensive** command is executed on an interface that is configured on a T4000 Type 5 FPC, the **IPv6 transit statistics** field displays:
 - Total statistics (sum of transit and local statistics) at the physical interface level
 - Transit statistics at the logical interface level
- When the **show interfaces extensive** command is executed on an aggregate interface in a T1600 Core Router, the IPv6 **Input bytes** is displayed for an aggregate interface. However, the IPv6 **Input bytes** is always zero on a member link of an aggregated bundle even when there is IPv6 transit traffic on the member link. This is because the logical interface index of the aggregate logical interface is updated but not the logical interface of the member links in the channel lookup table.
- The **Output packets** field under the **Traffic statistics** section in the output of the **show interfaces extensive** command includes both IPv4 and IPv6 packets. For example, in a scenario in which both IPv4 and IPv6 packets are being mirrored on the same interface and when you deactivate an IPv4 port-mirroring instance on the chassis, the output of the **show interfaces extensive** command shows a value in the Output packets field of the Traffic statistics section, which is the value of IPv6 packets that are mirrored and not of the IPv4 packets. This behavior is expected.
- For IQ2 PIC interfaces, the output of the **show interfaces extensive** command displays byte statistics that includes Layer 2 headers.
- If there are active OTN defects when an ISSU is performed, and the defect persists after the upgrade completes, the OTN alarm count is incremented by 1. For example, if an OTN alarm is active with a count of 1 and the defect remains after ISSU, the alarm count is incremented to 2. This behavior is expected.

This command has no options.

Required Privilege Level

view

List of Sample Output

[show interfaces extensive \(Circuit Emulation\) on page 992](#)

[show interfaces extensive \(Fast Ethernet\) on page 993](#)

[show interfaces extensive \(Gigabit Ethernet\) on page 996](#)

[show interfaces extensive \(10-Gigabit Ethernet\) on page 996](#)

[show interfaces extensive \(IQ2 and IQ2E\) on page 999](#)

[show interfaces extensive \(100-Gigabit Ethernet Type 4 PIC with CFP\) on page 1003](#)

[show interfaces extensive \(PTX5000 Packet Transport Router\) on page 1006](#)

[show interfaces extensive \(PTX10008 Routers\) on page 1009](#)

[show interfaces extensive \(PTX1000 Routers\) on page 1016](#)

[show interfaces extensive \(MX Series Routers\) on page 1018](#)

[show interfaces extensive \(MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface\) on page 1021](#)

[show interfaces extensive \(MX480 Router with MPC5E and 100-Gigabit Ethernet OTN Interface\) on page 1023](#)

[show interfaces extensive \(\(MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC\) on page 1027](#)

[show interfaces extensive \(PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC\) on page 1030](#)

[show interfaces extensive \(on page 1034](#)

[show interfaces extensive \(MX2020 Router with MPC6E and OTN MIC\) on page 1034](#)

[show interfaces extensive \(MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface\) on page 1038](#)

[show interfaces extensive \(MX2010 Router with MPC6E and 10-Gigabit Ethernet Interface\) on page 1040](#)

[show interfaces extensive \(T4000 Routers with Type 5 FPCs\) on page 1042](#)

[show interfaces extensive \(Aggregated Ethernet\) on page 1044](#)

Output Fields

For more information, see the output fields table for the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under [“Common Output Fields Description” on page 779](#). For sample output for specific interfaces, see the other topics in this collection.

Sample Output

show interfaces extensive (Circuit Emulation)

If a Circuit Emulation (CE) PIC is configured for SAToP pseudowire, then pseudowire statistics are displayed in the CE information section of the **show interface extensive** output. If SAToP pseudowire is not configured on the CE PIC, then all CE information counters display 0 (zero).

user@host> **show interface t1-0/0/0 extensive**

```
Physical interface :t1-0/0/0, Enabled, Physical Link : Up
  Interface index:61441
  Speed : 1.54 Mbps, Loopback: Disabled
  Operational state : Enabled,      Encapsulation : Trans
  Encoding : b8zs,      Framing      : unframe,      Build-out : 0-30
  Inversion      : enable, Clock source : master
  Description :
  Traffic statistics:
  T1 media:      Seconds
  ES              1643
  SES            1643

  CE Info      Packets      Bytes
  CE Rx       :    2395529    306627712
  CE Tx       :    2396259    306721152
  CE Rx Drop:      0          0
  CE Tx Drop:      0          0

  CE Overrun  Events: 0
  CE Underrun Events: 0
```

Sample Output

show interfaces extensive (Fast Ethernet)

user@host> **show interfaces fe-0/2/1 extensive**

```
Physical interface: fe-0/2/0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  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:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped      : 2006-04-16 23:00:41 PDT (02:08:05 ago)
  Statistics last cleared: 2006-04-16 21:42:00 PDT (03:26:46 ago)
  Traffic statistics:
    Input bytes      :          17539          152 bps
```

```

Output bytes      :          92968          224 bps
Input  packets:          348          0 pps
Output packets:         1349          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
Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort          66             66              0
  1 expedited-fo         0             0              0
  2 assured-forw         0             0              0
  3 network-cont        1283            1283             0
Active alarms  : None
Active defects : None
MAC statistics:
  Receive      Transmit
  Total octets    24721      105982
  Total packets   348        1349
  Unicast packets 347        430
  Broadcast packets 1         37
  Multicast packets 0         882
  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      348
  Input packet rejects    0
  Input DA rejects        0
  Input SA rejects        0
  Output packet count          1349
  Output packet pad count      0
  Output packet error count    0
  CAM destination filters: 3, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete

```

Link partner:

Link mode: Full-duplex, Flow control: None, Remote fault: OK

Packet Forwarding Engine configuration:

Destination slot: 0

CoS information:

CoS transmit queue		Bandwidth		Buffer	Priority	Limit
	%	bps	%	usec		
0 best-effort	95	95000000	95	0	low	none
3 network-control	5	5000000	5	0	low	none

Logical interface fe-0/2/0.0 (Index 66) (SNMP ifIndex 46) (Generation 133)

Flags: SNMP-Traps Encapsulation: ENET2

Protocol inet, MTU: 1500, Generation: 142, Route table: 0

Flags: DCU, SCU-out

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
silv1_new	0	0
(0)	(0)
silv2_new	0	0
(0)	(0)
silv_misc	0	0
(0)	(0)
silver0	0	0
(0)	(0)
silver2	0	0
(0)	(0)
silver3	0	0
(0)	(0)
silver4	0	0
(0)	(0)
silver5	0	0
(0)	(0)
silver6	0	0
(0)	(0)
silver7	0	0
(0)	(0)
silver9	0	0
(0)	(0)

Source class	Packets (packet-per-second)	Bytes (bits-per-second)
gold1	0	0
(0)	(0)
gold2	16600	1062400
(0)	(0)
gold3	0	0

```

(
0) (
0)
Addresses, Flags: Is-Preferred Is-Primary
Destination: 192.168.220.24/30, Local: 192.168.220.26, Broadcast:
192.168.220.27, Generation: 150

```

show interfaces extensive (Gigabit Ethernet)

```
user@host> show interfaces ge-5/0/0.0 extensive
```

```

Logical interface ge-5/0/0.0 (Index 71) (SNMP ifIndex 1930) (Generation 139)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Traffic statistics:
  Input  bytes   :                0
  Output bytes   :               42
  Input  packets :                0
  Output packets :                1
Local statistics:
  Input  bytes   :                0
  Output bytes   :               42
  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
Output Filters: f-any
Protocol inet, MTU: 1500, Generation: 155, Route table: 0
  Output Filters: f-inet,
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.168.220.24/30, Local: 192.168.220.26, Broadcast:
192.168.220.27,
    Generation: 170
Protocol multiservice, MTU: Unlimited, Generation: 156, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet)

```
user@host> show interfaces xe-2/1/0 extensive
```

```

Physical interface: xe-2/1/0, Enabled, Physical link is Up
  Interface index: 258, SNMP ifIndex: 762, Generation: 2046
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loopback: None, Source filtering: Disabled,
  Flow control: Enabled
  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:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped      : 2011-12-17 00:19:02 PST (07:36:37 ago)
  Statistics last cleared: 2011-12-17 07:55:24 PST (00:00:15 ago)
  Traffic statistics:
    Input bytes      :                110000                0 bps
    Output bytes     :                   0                0 bps
    Input packets    :                1000                0 pps
    Output packets   :                   0                0 pps
  IPv6 transit statistics:
    Input bytes      :                110000
    Output bytes     :                   0
    Input packets    :                1000
    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    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
PCS statistics
    Bit errors          0
    Errored blocks      0
MAC statistics:
    Receive             Transmit
    Total octets        128000      0
    Total packets       1000        0
    Unicast packets     1000        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   1000
    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: 2
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
Interface transmit statistics: Disabled

Logical interface xe-2/1/0.0 (Index 83) (SNMP ifIndex 1677) (Generation 10082)
Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
Traffic statistics:

```

```

Input  bytes :           110000
Output bytes :              0
Input  packets:           1000
Output packets:            0
IPv6 transit statistics:
  Input  bytes :           55000
  Output bytes :              0
  Input  packets:           500
  Output packets:            0
Local statistics:
  Input  bytes :           55000
  Output bytes :              0
  Input  packets:           500
  Output packets:            0
Transit statistics:
  Input  bytes :           55000          0 bps
  Output bytes :              0          0 bps
  Input  packets:           500          0 pps
  Output packets:            0          0 pps
IPv6 transit statistics:
  Input  bytes :           55000
  Output bytes :              0
  Input  packets:           500
  Output packets:            0
Protocol inet6, MTU: 1500, Generation: 23739, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 2001:0db8:0a0b:12f0:0000:0000:0000/112, Local:
2001:0db8:0a0b:12f0:0000:0000:0000:0001
  Generation: 506
  Addresses, Flags: Is-Preferred
    Destination: Odb8::/64, Local: Odb8::21d:b5ff:fef8:6deb
Protocol multiservice, MTU: Unlimited, Generation: 508
Generation: 23740, Route table: 0
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (IQ2 and IQ2E)

user@host> show interfaces ge-3/2/2 extensive

```

Physical interface: ge-3/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 548, Generation: 159
  Link-level type: Ethernet, MTU: 1518, 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
CoS queues       : 8 supported, 8 maximum usable queues
Schedulers       : 128
Hold-times       : Up 0 ms, Down 0 ms
Current address:  00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
Last flapped     : 2010-03-17 04:03:11 PDT (00:45:30 ago)
Statistics last cleared: Never
Traffic statistics:
  Input  bytes   :                1716096                0 bps
  Output bytes   :                1716448                0 bps
  Input  packets :                13407                0 pps
  Output packets :                13411                0 pps
IPv6 total statistics:
  Input  bytes   :                1716096
  Output bytes   :                1716096
  Input  packets :                13407
  Output packets :                13407
Ingress traffic statistics at Packet Forwarding Engine:
  Input  bytes   :                1716096                0 bps
  Input  packets :                13407                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: 1, 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: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

    0 best-effort                13407                13407
0
    1 expedited-fo                  0                  0
0
    2 assured-forw                  0                  0
0
    3 network-cont                  0                  0
0

Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

```



```

0      0 best-effort          13407          13407
0
0      1 expedited-fo         0            0
0
0      2 assured-forw         0            0
0
0      3 network-cont         4            4
0

Active alarms  : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets    13407        13411
Unicast packets  13407        13407
Broadcast packets      0            0
Multicast packets      0            4
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   13407
Input packet rejects  0
Input DA rejects     0
Input SA rejects     0
Output packet count           13411
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: None, Remote fault: OK
Local resolution:
Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 3
CoS information:
Direction : Output

```

Limit	CoS transmit queue	Bandwidth			Buffer	Priority
		%	bps	%		
0	best-effort	95	950000000	95	0	low
none						
3	network-control	5	50000000	5	0	low
none						
Direction : Input						
Limit	CoS transmit queue	Bandwidth			Buffer	Priority
		%	bps	%		
0	best-effort	95	950000000	95	0	low
none						
3	network-control	5	50000000	5	0	low
none						

Logical interface ge-3/2/2.0 (Index 83) (SNMP ifIndex 6080) (Generation 148)

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.100] Encapsulation: ENET2

Traffic statistics:

Input bytes : 0
Output bytes : 336
Input packets: 0
Output packets: 4

IPv6 total statistics:

Input bytes : 1716096
Output bytes : 1716096
Input packets: 13407
Output packets: 13407

Local statistics:

Input bytes : 0
Output bytes : 336
Input packets: 0
Output packets: 4

Transit statistics:

Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

IPv6 total statistics:

Input bytes : 1716096
Output bytes : 1716096
Input packets: 13407
Output packets: 13407

```

    Protocol inet6, MTU: 1500, Generation: 159, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Primary
        Destination: Unspecified, Local: 2000::2
    Generation: 146
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::214:f600:6412:86fa
    Protocol multiservice, MTU: Unlimited, Generation: 148
    Generation: 160, Route table: 0
      Policer: Input: __default_arp_policer__

Logical interface ge-3/2/2.32767 (Index 84) (SNMP ifIndex 6081) (Generation
149)
  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
  Protocol multiservice, MTU: Unlimited, Generation: 161, Route table: 0
    Flags: None
    Policer: Input: __default_arp_policer__

```

show interfaces extensive (100-Gigabit Ethernet Type 4 PIC with CFP)

```
user@host> show interfaces et-0/0/0:0 extensive
```

```

Physical interface: et-0/0/0:0, Enabled, Physical link is Down
  Interface index: 156, SNMP ifIndex: 516, Generation: 163
  Link-level type: Ethernet, MTU: 9192, Speed: 50000mbps, 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
CoS queues      : 8 supported, 8 maximum usable queues
Hold-times      : Up 0 ms, Down 0 ms
Damping         : half-life: 5 sec, max-suppress: 20 sec, reuse 1000, suppress:
2000, state: enabled
Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
Last flapped    : 2010-01-07 16:36:49 PST (18:02:35 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes      :                      0                      0 bps
Output bytes     :                      0                      0 bps
Input packets    :                      0                      0 pps
Output packets   :                      0                      0 pps
IPv6 transit statistics:
Input bytes      :                      0
Output bytes     :                      0
Input packets    :                      0
Output packets   :                      0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0,
L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors:
0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
0 DEFAULT, NC-        0                      0                      0
1 REALTIME            0                      0                      0
2 PRIVATE, NC-        0                      0                      0
3 CONTROL             1253                   1253                   0
4 BC-H, CLASS_        0                      0                      0
5 BC-M, CLASS_        0                      0                      0
6 IA, CLASS_V_        0                      0                      0
7 CLASS_S_OUTP        0                      0                      0
Queue      Mapped Forwarding Class
0          DEFAULT, NC-Q0
1          REALTIME
2          PRIVATE, NC-Q1
3          CONTROL
4          BC-H, CLASS-Q4
5          BC-M, CLASS-Q5

```

```

6          IA, CLASS_V_OUTPUT
7          CLASS_S_OUTPUT
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
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue            Bandwidth          Buffer Priority Limit
                                %          bps          %          usec
0 best-effort                   95      47500000000    95          0      low none
3 network-control                5      2500000000     5          0      low none

Logical interface et-0/0/0:0.0 (Index 68) (SNMP ifIndex 546) (Generation 161)
Flags: Deviet-Down SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes :                0
  Output bytes :                0
  Input packets:                0
  Output packets:               0
Local statistics:
  Input bytes :                0
  Output bytes :                0
  Input packets:                0
  Output packets:               0
Transit statistics:

```

```

      Input  bytes   :                0                0 bps
      Output bytes   :                0                0 bps
      Input  packets:                0                0 pps
      Output packets:                0                0 pps
Protocol inet, MTU: 9178, Generation: 220, Route table: 0
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 192.168.220.24/30, Local: 192.168.220.26, Broadcast:
192.168.220.27, Generation: 192
  Protocol mpls, MTU: 9166, Maximum labels: 3, Generation: 221, Route table: 0
  Protocol multiservice, MTU: Unlimited, Generation: 222, Route table: 0
  Policers: Input: __default_arp_policer

```

show interfaces extensive (PTX5000 Packet Transport Router)

user@host> show interfaces et-0/0/6 extensive

```

Physical interface: et-0/0/6, Enabled, Physical link is Up
  Interface index: 347, SNMP ifIndex: 531, Generation: 350
  Link-level type: Ethernet, MTU: 1514, Speed: 40Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Enabled
  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
  Damping            : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
state: unsuppressed
  Current address: 30:b6:4f:02:29:06, Hardware address: 30:b6:4f:02:29:06
  Last flapped      : 2017-02-15 21:40:06 PST (22:55:13 ago)
  Statistics last cleared: 2017-02-16 20:33:02 PST (00:02:17 ago)
  Traffic statistics:
    Input  bytes   :          1760000                0 bps
    Output bytes   :          1540000                0 bps
    Input  packets:           16000                0 pps
    Output packets:           14000                0 pps
  IPv6 transit statistics:
    Input  bytes   :          880000
    Output bytes   :          770000
    Input  packets:           8000
    Output packets:           7000
  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	7000	7000	0
1	0	0	0
2	0	0	0
3	7000	7000	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	2048000	1792000
Total packets	16000	14000
Unicast packets	16000	14000
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

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	38000000000	95	0	low
none					
3 network-control	5	2000000000	5	0	low
none					

Preclassifier statistics:

Traffic Class	Received Packets	Transmitted Packets	Dropped
best-effort	0	0	
0			
best-effort	0	0	
0			
best-effort	0	0	
0			
best-effort	0	0	
0			
best-effort	0	0	
0			
best-effort	0	0	
0			
best-effort	0	0	
0			

Link Degradation :

Link Monitoring : Disable

Interface transmit statistics: Disabled

Logical interface et-0/0/6.0 (Index 93) (SNMP ifIndex 841) (Generation 158)

Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2

Traffic statistics:

Input bytes : 1760000
 Output bytes : 1540000
 Input packets: 16000
 Output packets: 14000

IPv6 transit statistics:

Input bytes : 880000
 Output bytes : 770000


```

    Input  packets:                8000
    Output packets:                7000
Local statistics:
    Input  bytes   :                0
    Output bytes   :                0
    Input  packets:                0
    Output packets:                0
Transit statistics:
    Input  bytes   :       1760000      0 bps
    Output bytes   :       1540000      0 bps
    Input  packets:        16000      0 pps
    Output packets:        14000      0 pps
IPv6 transit statistics:
    Input  bytes   :       880000
    Output bytes   :       770000
    Input  packets:        8000
    Output packets:        7000
Protocol inet, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
  Generation: 206, Route table: 0
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 11.0.0/24, Local: 11.0.0.2, Broadcast: 11.0.0.255, Generation:
228
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
      Generation: 207, Route table: 0
        Addresses, Flags: Is-Preferred Is-Primary
          Destination: 1100::/120, Local: 1100::2
      Generation: 230
        Addresses, Flags: Is-Preferred
          Destination: fe80::/64, Local: fe80::32b6:4fff:fe02:2906
    Protocol multiservice, MTU: Unlimited, Generation: 232
    Generation: 208, Route table: 0
      Policer: Input: __default_arp_policer__

```

show interfaces extensive (PTX10008 Routers)

```
user@host> show interfaces ae0 extensive
```

```

Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 917, SNMP ifIndex: 817, Generation: 4436
  Link-level type: Ethernet, MTU: 1518, Speed: 20Gbps, BPDU Error: None, MAC-REWRITE
  Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control:
  Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 30:b6:4f:e9:7c:05, Hardware address: 30:b6:4f:e9:7c:05
  Last flapped   : 2017-04-10 05:20:29 PDT (00:03:52 ago)
  Statistics last cleared: 2017-04-10 05:21:52 PDT (00:02:29 ago)
  Traffic statistics:
    Input bytes  :          36463816334          0 bps
    Output bytes :          36463816334          0 bps
    Input packets:          24671053          0 pps
    Output packets:          24671053          0 pps
  IPv6 transit statistics:
    Input bytes  :          18231905950
    Output bytes :          18231905950
    Input packets:          12335525
    Output packets:          12335525
  MAC statistics:
    Broadcast packets          Receive          Transmit
    Multicast packets          0              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
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0                24671053            24671053          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 ae0.0 (Index 99) (SNMP ifIndex 832) (Generation 43813)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

```

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

Bundle:

Input :	4934211	0	7292763858	0
Output:	4934211	0	7292763858	0

Adaptive Statistics:

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

Link:

et-0/0/28:0.0

Input :	4934211	0	7292763858	0
Output:	4934211	0	7292763858	0

et-0/0/28:3.0

Input :	0	0	0	0
Output:	0	0	0	0

Aggregate member links: 2

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

et-0/0/28:0.0	0	0	0	0
---------------	---	---	---	---

et-0/0/28:3.0	0	0	0	0
---------------	---	---	---	---

Protocol inet, MTU: 1500

Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold cnt: 0, NH drop cnt: 0

Generation: 89219, Route table: 0

Flags: Sendbroadcast-pkt-to-re

Addresses, Flags: Is-Preferred Is-Primary

Destination: 21.0.0.0/30, Local: 21.0.0.1, Broadcast: 21.0.0.3, Generation: 62420

Protocol inet6, MTU: 1500

Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold cnt: 0, NH drop cnt: 0

Generation: 89220, Route table: 0

Addresses, Flags: Is-Preferred Is-Primary

Destination: 3001::1500:0/126, Local: 3001::1500:1

Generation: 62422

Addresses, Flags: Is-Preferred

Destination: fe80::/64, Local: fe80::32b6:4f00:2e9:7c05

Protocol multiservice, MTU: Unlimited, Generation: 62424

Generation: 89221, Route table: 0

Policer: Input: __default_arp_policer__

Logical interface ae0.1 (Index 100) (SNMP ifIndex 833) (Generation 43814)

```

Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.3 ] Encapsulation: ENET2
Statistics          Packets          pps          Bytes          bps
Bundle:
  Input :           4934211           0    7292763858           0
  Output:           4934211           0    7292763858           0
Adaptive Statistics:
  Adaptive Adjusts:           0
  Adaptive Scans  :           0
  Adaptive Updates:           0
Link:
  et-0/0/28:0.1
    Input :           0           0           0           0
    Output:           4934211           0    7292763858           0
  et-0/0/28:3.1
    Input :           4934211           0    7292763858           0
    Output:           0           0           0           0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
  et-0/0/28:0.1           0           0           0           0
  et-0/0/28:3.1           0           0           0           0
Protocol inet, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 89222, Route table: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 21.0.0.4/30, Local: 21.0.0.5, Broadcast: 21.0.0.7, Generation:
62426
Protocol inet6, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 89223, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 3001::1500:4/126, Local: 3001::1500:5
Generation: 62428
Addresses, Flags: Is-Preferred
Destination: fe80::/64, Local: fe80::32b6:4f00:3e9:7c05
Protocol multiservice, MTU: Unlimited, Generation: 62430
Generation: 89224, Route table: 0
Policer: Input: __default_arp_policer__

Logical interface ae0.2 (Index 101) (SNMP ifIndex 834) (Generation 43815)
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4 ] Encapsulation: ENET2
Statistics          Packets          pps          Bytes          bps
Bundle:

```

```

    Input :          4934211          0    7292763858          0
    Output:          4934211          0    7292763858          0
Adaptive Statistics:
    Adaptive Adjusts:          0
    Adaptive Scans :          0
    Adaptive Updates:          0
Link:
    et-0/0/28:0.2
        Input :          2467106          0    3646382668          0
        Output:          4934211          0    7292763858          0
    et-0/0/28:3.2
        Input :          2467105          0    3646381190          0
        Output:              0          0          0          0
Marker Statistics:   Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
    et-0/0/28:0.2              0              0              0              0
    et-0/0/28:3.2              0              0              0              0
Protocol inet, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 89225, Route table: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
    Destination: 21.0.0.8/30, Local: 21.0.0.9, Broadcast: 21.0.0.11, Generation:
62432
Protocol inet6, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 89226, Route table: 0
Addresses, Flags: Is-Preferred Is-Primary
    Destination: 3001::1500:8/126, Local: 3001::1500:9
Generation: 62434
Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::32b6:4f00:4e9:7c05
Protocol multiservice, MTU: Unlimited, Generation: 62436
Generation: 89227, Route table: 0
Policer: Input: __default_arp_policer__

Logical interface ae0.3 (Index 102) (SNMP ifIndex 835) (Generation 43816)
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.5 ] Encapsulation: ENET2
Statistics          Packets          pps          Bytes          bps
Bundle:
    Input :          4934210          0    7292762380          0
    Output:          4934210          0    7292762380          0
Adaptive Statistics:

```

```

    Adaptive Adjusts:      0
    Adaptive Scans   :      0
    Adaptive Updates:      0
Link:
  et-0/0/28:0.3
    Input :      4934210      0      7292762380      0
    Output:      0      0      0      0
  et-0/0/28:3.3
    Input :      0      0      0      0
    Output:      4934210      0      7292762380      0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
  et-0/0/28:0.3      0      0      0      0
  et-0/0/28:3.3      0      0      0      0
Protocol inet, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
  Generation: 89228, Route table: 0
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 21.0.0.12/30, Local: 21.0.0.13, Broadcast: 21.0.0.15,
Generation: 62438
  Protocol inet6, MTU: 1500
  Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold
cnt: 0, NH drop cnt: 0
  Generation: 89229, Route table: 0
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 3001::1500:c/126, Local: 3001::1500:d
Generation: 62440
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::32b6:4f00:5e9:7c05
Protocol multiservice, MTU: Unlimited, Generation: 62442
Generation: 89230, Route table: 0
  Policer: Input: __default_arp_policer__

Logical interface ae0.4 (Index 103) (SNMP ifIndex 836) (Generation 43817)
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.6 ] Encapsulation: ENET2
  Statistics      Packets      pps      Bytes      bps
  Bundle:
    Input :      4934210      0      7292762380      0
    Output:      4934210      0      7292762380      0
  Adaptive Statistics:
    Adaptive Adjusts:      0
    Adaptive Scans   :      0
    Adaptive Updates:      0

```

```

Link:
  et-0/0/28:0.4
    Input :      2467105      0      3646381190      0
    Output:      2467105      0      3646381190      0
  et-0/0/28:3.4
    Input :      2467105      0      3646381190      0
    Output:      2467105      0      3646381190      0
Marker Statistics:  Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
  et-0/0/28:0.4              0              0              0              0
  et-0/0/28:3.4              0              0              0              0
Protocol inet, MTU: 1500
Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
Generation: 89231, Route table: 0
  Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 21.0.0.16/30, Local: 21.0.0.17, Broadcast: 21.0.0.19,
Generation: 62444
  Protocol inet6, MTU: 1500
  Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 2, Curr new hold
cnt: 0, NH drop cnt: 0
  Generation: 89232, Route table: 0
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 3001::1500:10/126, Local: 3001::1500:11
  Generation: 62446
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::32b6:4f00:6e9:7c05
  Protocol multiservice, MTU: Unlimited, Generation: 62448
  Generation: 89233, Route table: 0
    Policer: Input: __default_arp_policer__

Logical interface ae0.32767 (Index 104) (SNMP ifIndex 5645) (Generation 43818)
  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:
    et-0/0/28:0.32767
      Input :            0            0            0            0

```

```

      Output:          0          0          0          0
et-0/0/28:3.32767
      Input  :          0          0          0          0
      Output:          0          0          0          0
Marker Statistics:   Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
et-0/0/28:0.32767           0          0          0          0
et-0/0/28:3.32767           0          0          0          0
Protocol multiservice, MTU: Unlimited, Generation: 89234, Route table: 0
Flags: None
Policer: Input: __default_arp_policer__

```

show interfaces extensive (PTX1000 Routers)

user@host> show interfaces et-0/0/48:1 extensive

```

Physical interface: et-0/0/48:1, Enabled, Physical link is Up
  Interface index: 306, SNMP ifIndex: 697, Generation: 311
  Link-level type: Ethernet, MTU: 1514, 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
  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
  Damping           : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
state: unsuppressed
  Current address: ec:13:db:62:4a:f6, Hardware address: ec:13:db:62:4a:f6
  Last flapped     : 2017-05-08 11:07:59 PDT (12:08:13 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes      :          0          0 bps
    Output bytes     :          0          0 bps
    Input packets    :          0          0 pps
    Output packets   :          0          0 pps
  IPv6 transit statistics:
    Input bytes      :          0
    Output bytes     :          0
    Input packets    :          0
    Output packets   :          0
  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

Egress queues: 8 supported, 4 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0	0	0	0
1	0	0	0
2	0	0	0
3	0	0	0

Queue number:	Mapped forwarding classes
0	best-effort
1	expedited-forwarding
2	assured-forwarding
3	network-control

Active alarms : None

Active defects : None

PCS statistics	Seconds
Bit errors	3
Errored blocks	3

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

```

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
  Link Degradate :
    Link Monitoring           : Disable
  Interface transmit statistics: Disabled

```

show interfaces extensive (MX Series Routers)

user@host> show interfaces xe-0/0/0 extensive

```

Physical interface: xe-0/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 592, Generation: 148
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2013-10-26 03:20:40 test (2w3d 03:15 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: 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                    0                0                        0
1                    0                0                        0
2                    0                0                        0
3                    0                0                        0
Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control
Active alarms : LINK
Active defects : LINK
PCS statistics      Seconds
  Bit errors        109
  Errored blocks    109
MAC statistics:      Receive      Transmit
1                    0            0            0
2                    0            0            0
3                    0            0            0
Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control
Active alarms : LINK
Active defects : LINK
PCS statistics      Seconds
  Bit errors        109
  Errored blocks    109
MAC statistics:      Receive      Transmit
  Total octets      0            0

```

```

Total packets                0                0
Unicast packets              0                0
Broadcast packets            0                0
Multicast packets            0                0
CRC/Align errors             0                0
FIFO errors                  0                0
MAC control frames           0                0
MAC pause frames             0                0
Oversized frames             0
Jabber frames                0
Fragment frames              0
VLAN tagged frames           0
Code violations              0
Total errors                  0                0
Filter statistics:
  Input packet count          0
  Input packet rejects        0
  Input DA rejects            0
  Input SA rejects            0
  Output packet count         0
  Output packet pad count     0
  Output packet error count   0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
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
  Interface transmit statistics: Disabled

```

When an ASIC is wedged, the interfaces are brought down along with the IFD. The reason for the link down is displayed as **ASIC-Error** in the **Device flags**.

user@host> **show interfaces xe-1/0/0 extensive**

```

Physical interface: xe-1/0/0, Administratively down, Physical link is Down
Interface index: 147, SNMP ifIndex: 563, Generation: 150
Link-level type: Ethernet, MTU: 1514, MRU: 0, LAN-PHY mode, Speed: 10Gbps, BPDU

```

```

Error: None, Loop Detect PDU Error: None,
MAC-REWRITE Error: None, Loopback: None, Source filtering: Disabled, Flow control:
Disabled
Pad to minimum frame size: Disabled
Device flags      : Present Running Down ASIC-Error
Interface flags: Hardware-Down Down 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: cc:e1:7f:a8:05:4a, Hardware address: cc:e1:7f:a8:05:4a
Last flapped      : 2017-06-05 17:20:54 PDT (00:03:51 ago)
Statistics last cleared: Never

```

show interfaces extensive (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

user@host> show interfaces xe-0/0/3 extensive

```

Physical interface: xe-0/0/3, Enabled, Physical link is Up
Interface index: 200, SNMP ifIndex: 577, Generation: 203
Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: None, 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
Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
Last flapped      : 2014-06-26 18:16:50 PDT (04:58:35 ago)
Statistics last cleared: Never
Traffic statistics:
Input  bytes   :                0                0 bps
Output bytes   :                0                0 bps
Input  packets :                0                0 pps
Output packets :                0                0 pps
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: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
    Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
    0                  0                0                      0
    1                  0                0                      0
    2                  0                0                      0
    3                  0                0                      0
Queue number:      Mapped forwarding classes
    0               best-effort
    1               expedited-forwarding
    2               assured-forwarding
    3               network-control
Active alarms  : None
Active defects : None
PCS statistics          Seconds
    Bit errors          0
    Errored blocks      4
MAC statistics:      Receive      Transmit
    Total octets        0          0
    Total packets       0          0
    Unicast packets     0          0
    Broadcast packets   0          0
    Multicast packets   0          0
    CRC/Align errors    0          0
    FIFO errors         0          0
    MAC control frames  0          0
    MAC pause frames    0          0
    Oversized frames    0
    Jabber frames       0
    Fragment frames     0
    VLAN tagged frames  0
    Code violations     0

```

```

    Total errors                                0                0
Filter statistics:
    Input packet count                        0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                      0
    Output packet pad count                  0
    Output packet error count                0
    CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
    Destination slot: 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
    Interface transmit statistics: Disabled

```

show interfaces extensive (MX480 Router with MPC5E and 100-Gigabit Ethernet OTN Interface)

user@host> show interfaces et-2/1/0 extensive

```

Physical interface: et-2/1/0, Enabled, Physical link is Up
  Interface index: 215, SNMP ifIndex: 872, Generation: 218
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU 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
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped    : 2014-06-26 18:42:04 PDT (04:36:58 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: 263, Errors: 0, Drops: 0, Collisions: 0, Aged packets:
0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0                      0                0                      0
  1                      0                0                      0
  2                      0                0                      0
  3                      0                0                      0
Queue number:      Mapped forwarding classes
  0                  best-effort
  1                  expedited-forwarding
  2                  assured-forwarding
  3                  network-control
Active alarms  : None
Active defects : None
PCS statistics                Seconds
  Bit errors                0
  Errored blocks            754
MAC statistics:                Receive      Transmit
  Total octets              14960          0
  Total packets              104           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

```



```

    0 best-effort          95    95000000000    95          0    low
none
    3 network-control      5     5000000000    5          0    low
none
    Interface transmit statistics: Disabled

```

show interfaces extensive ((MX960 Router with MPC3E and 100-Gigabit DWDM OTN MIC)

user@host> show interfaces et-3/0/0 extensive

```

Physical interface: et-3/0/0, Enabled, Physical link is Up
  Interface index: 163, SNMP ifIndex: 564, Generation: 166
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU 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
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  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: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2016-02-17 14:26:31 PST (09:04:28 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
  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: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU
errors: 0,
  Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0                      0                0                      0
  1                      0                0                      0
  2                      0                0                      0
  3                      0                0                      0
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms   : None
Active defects  : None
PCS statistics          Seconds
  Bit errors              8
  Errored blocks          10
MAC statistics:      Receive      Transmit
  Total octets              0          0
  Total packets              0          0
  Unicast packets            0          0
  Broadcast packets          0          0
  Multicast packets          0          0
  CRC/Align errors          0          0
  FIFO errors                0          0
  MAC control frames         0          0
  MAC pause frames           0          0
  Oversized frames           0
  Jabber frames              0
  Fragment frames            0
  VLAN tagged frames         0
  Code violations            0
  Total errors               0          0
Filter statistics:
  Input packet count         0
  Input packet rejects       0

```

```

Input DA rejects          0
Input SA rejects          0
Output packet count              0
Output packet pad count         0
Output packet error count       0
CAM destination filters: 0, CAM source filters: 0
OTN alarms      : None
OTN defects     : None
OTN FEC Mode    : SDFEC
OTN Rate       :      OTU4 (120.5Gbps)
OTN Line Loopback : None
OTN Local Loopback: None
OTN Payload PRBS  : None
OTN Laser Enable  : On
OTN FEC statistics:
    Corrected Errors          7065332638
    Uncorrected Words         3412572
    Corrected Error Ratio (    32785 sec average) 1.79e-06 (INVALID)
OTN FEC alarms:      Seconds      Count  State
    FEC Degrade          0          0  OK
    FEC Excessive        3          1  OK
OTN OC:              Seconds      Count  State
    LOS                   3          1  OK
    LOF                   50         3  OK
    LOM                   3          3  OK
    Wavelength Lock       0          0  OK
OTN OTU:
    AIS                   0          0  OK
    BDI                   4          4  OK
    IAE                   4          4  OK
    TTIM                  1          1  OK
    BIAE                  3          3  OK
    TSF                   50         3  OK
    SSF                   50         3  OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN ODU:
    AIS                   20          2  OK
```

```

OCI                      4          4 OK
LCK                      4          4 OK
BDI                      2          2 OK
TTIM                    20          2 OK
IAE                     0          0 OK
LTC                     0          0 OK
CSF                     18          2 OK
TSF                     66          2 OK
SSF                     66          2 OK
PTIM                    43          2 OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x07
ODU Delay Management :
  Result   : 0ms
PRBS:
  Result: Test not enabled
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x07
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      95000000000      95           0      low
none
    3 network-control         5       50000000000       5           0      low
none
  Interface transmit statistics: Disabled

```

show interfaces extensive (PTX3000 Router with 5-port 100-Gigabit DWDM OTN PIC)

user@host > **show interfaces extensive et-4/0/0**

```

Physical interface: et-4/0/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 544, Generation: 161
  Link-level type: Ethernet, MTU: 1514, Speed: 100Gbps, BPDU Error: None, Loopback:
Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags      : Present Running
  Interface flags: 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
  Damping          : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0,
state: unsuppressed
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped    : 2016-06-04 21:42:42 PDT (1d 05:09 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: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0                0                0                0
    1                0                0                0
    2                0                0                0
    3                0                0                0
  Queue number:
    Mapped forwarding classes
    0                best-effort
    1                expedited-forwarding
    2                assured-forwarding

```

```

3                      network-control
Active alarms   : None
Active defects  : None
PCS statistics          Seconds
  Bit errors              7
Errored blocks        10
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
OTN alarms       : None
OTN defects      : None
OTN FEC Mode     : SDFEC
OTN Rate         :      OTU4 (120.5Gbps)
OTN Line Loopback : None
OTN Local Loopback: None
OTN Payload PRBS  : None
OTN Laser Enable  : On
OTN FEC statistics:
  Corrected Errors              19637746
  Uncorrected Words              0
  Corrected Error Ratio (      104923 sec average)  1.55e-09
OTN FEC alarms:      Seconds      Count  State
  FEC Degrade         0          0  OK

```


FEC Excessive	0	0	OK
---------------	---	---	----

OTN OC:	Seconds	Count	State
LOS	0	0	OK
LOF	2	1	OK
LOM	2	1	OK
Wavelength Lock	0	0	OK

OTN OTU:

AIS	0	0	OK
BDI	2	1	OK
IAE	0	0	OK
TTIM	0	0	OK
BIAE	0	0	OK
TSF	2	1	OK
SSF	0	0	OK

Received DAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Received SAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted DAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted SAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
OTN ODU:

AIS	0	0	OK
OCI	0	0	OK
LCK	0	0	OK
BDI	2	1	OK
TTIM	0	0	OK
IAE	0	0	OK
LTC	0	0	OK
CSF	0	0	OK
TSF	2	1	OK
SSF	0	0	OK
PTIM	2	1	OK

Received DAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Received SAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted DAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Transmitted SAPI:

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
OTN Received Overhead Bytes:

APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00

```

    Payload Type: 0x07
    ODU Delay Management :
        Result : 0ms
    PRBS:
        Result: Test not enabled
    OTN Transmitted Overhead Bytes:
        APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x07
    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      950000000000    95          0      low
    none
        3 network-control          5       50000000000     5          0      low
    none
    Interface transmit statistics: Disabled

```

show interfaces extensive (

show interfaces extensive (MX2020 Router with MPC6E and OTN MIC)

user@host> **show interfaces xe-3/0/0 extensive**

```

Physical interface: xe-3/0/0, Enabled, Physical link is Up
    Interface index: 166, SNMP ifIndex: 516, Generation: 169
    Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
    BPDU Error:
    None, MAC-REWRITE Error: None, Loopback: None, 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
        Hold-times : Up 0 ms, Down 0 ms
        Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
        Last flapped : 2014-05-28 17:53:12 PDT (05:56:24 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: 3, 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-forwarding      0          0
  0
  2 assured-forwarding      0          0
  0
  3 network-control      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      2
  Errored blocks      2
MAC statistics:      Receive      Transmit
  Total octets      0      0

```

Total packets	0	0
Unicast packets	0	0
Broadcast packets	0	0
Multicast packets	0	0
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	
Total errors	0	0
Filter statistics:		
Input packet count	0	
Input packet rejects	0	
Input DA rejects	0	
Input SA rejects	0	
Output packet count		0
Output packet pad count		0
Output packet error count		0
CAM destination filters: 0, CAM source filters: 0		
OTN alarms	: None	
OTN defects	: None	
OTN FEC Mode	: GFEC	
OTN Rate	: Fixed Stuff Bytes 11.0957Gbps	
OTN Line Loopback	: None	
OTN Local Loopback	: None	
OTN Payload PRBS	: None	
OTN FEC statistics:		
Corrected Errors		0
Uncorrected Words		0
Corrected Error Ratio (21387 sec average)	0.00e+00
OTN FEC alarms:	Seconds	Count State
FEC Degrade	0	0 OK
FEC Excessive	0	0 OK
OTN OC:	Seconds	Count State
LOS	0	0 OK
LOF	0	0 OK
LOM	0	0 OK
Wavelength Lock	0	0 OK
OTN OTU:		
AIS	0	0 OK

```

BDI                                0                0 OK
IAE                                0                0 OK
TTIM                               0                0 OK
BIAE                               0                0 OK
TSF                                0                0 OK
SSF                                0                0 OK
Received DAPI:
00 53 4d 2d 54 52 43 20 44 41 50 49 2d 53 45 43 .SM-TRC DAPI-SEC
Received SAPI:
00 53 4d 2d 54 52 43 20 53 41 50 49 2d 53 45 43 .SM-TRC SAPI-SEC
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN ODU:
AIS                                0                0 OK
OCI                                0                0 OK
LCK                                0                0 OK
BDI                                0                0 OK
TTIM                               0                0 OK
IAE                                0                0 OK
LTC                                0                0 OK
CSF                                0                0 OK
TSF                                0                0 OK
SSF                                0                0 OK
PTIM                               0                0 OK
Received DAPI:
00 50 4d 2d 54 52 43 20 44 41 50 49 2d 53 45 43 .PM-TRC DAPI-SEC
Received SAPI:
00 50 4d 2d 54 52 43 20 53 41 50 49 2d 53 45 43 .PM-TRC SAPI-SEC
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x00
ODU Delay Management :
Result : 0x00
PRBS:
Result: Test not enabled
OTN Transmitted Overhead Bytes:
APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x00

```

```

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
  Interface transmit statistics: Disabled

```

show interfaces extensive (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

user@host> show interfaces et-9/0/0 extensive

```

Physical interface: et-9/0/0, Enabled, Physical link is Up
  Interface index: 196, SNMP ifIndex: 623, Generation: 199
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, Speed: 100Gbps, BPDU 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
  Hold-times        : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped      : 2014-06-26 18:18:34 PDT (04:17:07 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                    0                0                        0
1                    0                0                        0
2                    0                0                        0
3                    0                0                        0
Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control
Active alarms   : None
Active defects  : None
PCS statistics          Seconds
  Bit errors              0
  Errored blocks          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
Total errors          0          0
Filter statistics:
  Input packet count    0
  Input packet rejects  0

```

```

Input DA rejects          0
Input SA rejects          0
Output packet count              0
Output packet pad count        0
Output packet error count      0
CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth          Buffer Priority
Limit
                                %          bps          %          usec
    0 best-effort              95    95000000000    95          0      low
none
    3 network-control          5     50000000000     5          0      low
none
Interface transmit statistics: Disabled

```

show interfaces extensive (MX2010 Router with MPC6E and 10-Gigabit Ethernet Interface)

user@host> show interfaces xe-6/1/0 extensive

```

Physical interface: xe-6/1/0, Enabled, Physical link is Up
  Interface index: 159, SNMP ifIndex: 603, Generation: 162
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: None, 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
  Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
  Last flapped   : 2014-06-26 18:16:50 PDT (04:21:04 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                    0                0                    0
  1                    0                0                    0
  2                    0                0                    0
  3                    0                0                    0
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms   : None
Active defects  : None
PCS statistics          Seconds
  Bit errors          0
  Errored blocks      1
MAC statistics:      Receive      Transmit
  Total octets        0            0
  Total packets       0            0
  Unicast packets     0            0
  Broadcast packets   0            0
  Multicast packets   0            0
  CRC/Align errors    0            0
  FIFO errors         0            0
MAC control frames    0            0
  MAC pause frames    0            0
  Oversized frames    0

```

```

Jabber frames                0
Fragment frames              0
VLAN tagged frames           0
Code violations               0
Total errors                  0          0
Filter statistics:
  Input packet count          0
  Input packet rejects        0
  Input DA rejects            0
  Input SA rejects            0
  Output packet count         0
  Output packet pad count     0
  Output packet error count   0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 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
  Interface transmit statistics: Disabled

```

show interfaces extensive (T4000 Routers with Type 5 FPCs)

The output fields for the **show interfaces *interface* extensive** command remains the same for 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP), 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP), and 100-Gigabit Ethernet Type 5 PIC with CFP (PF-1CGE-CFP).

user@host> **show interfaces xe-4/0/0 extensive**

```

Physical interface: xe-4/0/0, Enabled, Physical link is Up
  Interface index: 200, SNMP ifIndex: 592, Generation: 203
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error:
None, Loopback: None, Source filtering: Disabled, Flow control: Enabled
  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

```

```

Damping      : half-life: 5 sec, max-suppress: 20 sec, reuse 1000, suppress:
2000, state: enabled
Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
Last flapped  : 2013-06-03 16:01:56 PDT (06:04:07 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: 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
PCS statistics      Seconds
  Bit errors              0
  Errored blocks          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
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    0                0                0
best-effort        0                0                0
Interface transmit statistics: Disabled

```

show interfaces extensive (Aggregated Ethernet)

user@host> show interfaces ae0 extensive

```

Physical interface: ae0, Enabled, Physical link is Up
Interface index: 199, SNMP ifIndex: 570, Generation: 202
Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, 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: SNMP-Traps Internal: 0x4000
Current address: 00:00:5E:00:53:00, Hardware address: 00:00:5E:00:53:00
Last flapped   : 2012-06-06 23:33:03 PDT (00:00:58 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes   :          18532          1984 bps
  Output bytes  :           0         0 bps
  Input packets:          158         2 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, 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 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: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort      57             57             0
  1 expedited-fo     0              0              0
  2 assured-forw     0              0              0
  3 network-cont    63605          63605          0
Queue number:        Mapped forwarding classes
  0                  best-effort
  1                  expedited-forwarding
  2                  assured-forwarding
  3                  network-control

```

```

Logical interface ae0.0 (Index 331) (SNMP ifIndex 583) (Generation 142)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  Statistics          Packets          pps          Bytes          bps
Bundle:
  Input :             149             2          17416          1984
  Output:              0             0              0              0
Link:
  ge-3/2/5.0
    Input :             90             1          10100          992
    Output:              0             0              0              0
  ge-3/3/9.0
    Input :             59             1           7316          992
    Output:              0             0              0              0
LACP info:          Role          System          System          Port
Port  Port
                                priority          identifier  priority          number
key
ge-3/2/5.0    Actor          100    00:00:00:00:00:01          127          1
1
ge-3/2/5.0    Partner          127    00:24:dc:98:67:c0          127          1    1
ge-3/3/9.0    Actor          100    00:00:00:00:00:01          127          2
1
ge-3/3/9.0    Partner          127    00:24:dc:98:67:c0          127          2    1
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
ge-3/2/5.0            38          137              0              0
ge-3/3/9.0            36          139              0              0
Marker Statistics:    Marker Rx      Resp Tx      Unknown Rx      Illegal Rx
ge-3/2/5.0              0              0              0              0
ge-3/3/9.0              0              0              0              0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
  Flags: Sendbcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255, Generation:
153
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer__

```

show interfaces fabric

Syntax

```
show interfaces fabric
<interface-name>
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced in Junos OS Release 12.3 for the QFX Series.

Description

Display status information about the specified fabric interface.

Options

interface-name—(QFabric systems only) Either the serial number or the alias of the QFabric switch component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and not contain any colons.

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.

routing-instance (all | instance-name)—(Optional) Display all routing instances or the name of an individual routing instance.

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

[Monitoring Interface Status and Traffic](#) | 392

List of Sample Output

[show interfaces fabric on page 1057](#)

[show interfaces fabric brief on page 1057](#)

[show interfaces fabric detail on page 1069](#)

[show interfaces fabric extensive on page 1070](#)

[show interfaces fabric terse on page 1073](#)

[show interfaces fabric device-name on page 1073](#)

Output Fields

[Table 83 on page 1048](#) lists the output fields for the **show interfaces fabric** command. Output fields are listed in the approximate order in which they appear.

Table 83: show interfaces fabric Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Type	Physical interface type; for example, Ethernet.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Clocking	Reference clock source.	detail
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex.	All levels

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC-REWRITE Error	Specifies if the encapsulation of the packet has been changed.	none
BPDU Error	Specifies if a BPDU has been received on a blocked interface.	none
Loopback	Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.	All levels
Source filtering	Source filtering status: Enabled or Disabled.	All levels
Flow control	Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Hold-Times	Current interface hold-time up and hold-time down, in milliseconds.	detail
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) .	detail extensive none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is Statistics last cleared: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, 2010-05-17 07:51:28 PDT (00:04:33 ago) .	detail extensive

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail extensive
IPv6 transit statistics	<p>If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail extensive

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runt—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output rate	Output rate in bps and pps.	None specified
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of packets that exceed the configured MTU. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet Forwarding Engine Configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Item	Type of QFabric system component being viewed. Possible values include Node group, Interconnect device, Fabric control, Fabric manager, Diagnostic routing engine, and Ungrouped Node device.	none
Identifier	Hardware serial identifier of a QFabric system component. When you configure an alias name for a component, the ID is displayed.	none
Connection	Status of a QFabric system component: either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for the listed component.	none
Configuration	Whether or not the configuration for a QFabric system component has been received and installed. The configuration can be Configured, Failed (unsuccessful), Pending (in the process of being written or retried), or Unknown.	none
Node group	Name of the Node groups associated with the QFabric system, and the Node devices assigned to each Node group. The group can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for the devices in the group. This field also displays the serial ID for the Node group and the status for the Node group.	none

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Fabric control	Name of the virtual Junos Routing Engines responsible for route selection within a QFabric system partition. The fabric control Routing Engine can be either Connected or Disconnected, depending on whether or not the Director software has detected keepalive messages for this virtual device. It also displays the identifier and configuration status for the fabric control Routing Engine.	none
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Flags	Information about the logical interface. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	All levels
Encapsulation	Encapsulation method used on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. NOTE: The bandwidth bps counter is not enabled.	detail extensive
Local statistics	<ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. NOTE: The bandwidth bps counter is not enabled.	detail extensive

Table 83: show interfaces fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	<ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail extensive
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet, the IP address of the interface is also displayed.	brief
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none

Sample Output

show interfaces fabric

```
user@switch> show interfaces fabric
```

Item	Identifier	Connection	Configuration
Node group			
	BBAK3775	Connected	Configured
	NW-NG-0	Connected	Configured
	P2659-C	Connected	Configured
	ptor-0	Connected	Configured
Fabric control			
	FC-0	Connected	Configured
	FC-1	Connected	Configured

show interfaces fabric brief

```
user@switch> show interfaces fabric brief
```

```
Physical interface: BBAK0372:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
```

```
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x0
```

```
Logical interface BBAK0372:fte-0/1/0.32768
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  eth-switch
```

```
Physical interface: BBAK0372:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
```

```
Logical interface BBAK0372:fte-0/1/2.32768
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  eth-switch
```

```
Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
```

```
Logical interface BBAK0394:fte-0/1/0.32768
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch
```

```
Physical interface: BBAK0394:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
```

```
Logical interface BBAK0394:fte-0/1/2.32768
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch
```

```
Physical interface: BBAK3775:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
Speed: Unspecified
  Device flags      : Present Running
```

```
Logical interface BBAK3775:bme0.0
  Flags: LinkAddress 0-0 Encapsulation: ENET2
```

```

inet  128.0.0.1/2
      128.0.0.16/2
      128.0.32.0/2
tnp    0x10

```

```

Logical interface BBAK3775:bme0.1
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  inet  128.0.0.13/2
        128.0.130.0/2

```

```

Logical interface BBAK3775:bme0.2
  Flags: Encapsulation: ENET2
  inet  128.0.0.13/8
        128.0.130.0/8
        169.254.128.13/16
        169.254.193.0/16

```

```

Physical interface: BBAK3775:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified
  Device flags      : Present Running
  Interface flags: SNMP-Traps

```

```

Logical interface BBAK3775:qfabric.0
  Flags: SNMP-Traps Encapsulation: ENET2
  inet
  mpls
  eth-switch

```

```

Physical interface: BBAK3775:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:
1000mbps
  Device flags      : Present Running

```

```

Logical interface BBAK3775:vcp0.32769
  Flags: LinkAddress 0-0 Encapsulation: ENET2

```

```

Physical interface: BBAK3775:vcpl, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
  Device flags      : Present Running

```

```

Logical interface BBAK3775:vcpl.32768
  Flags: LinkAddress 0-0 Encapsulation: ENET2

```

```

Physical interface: BBAK3775:vcp2, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
  Device flags    : Present Running

  Logical interface BBAK3775:vcp2.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: BBAK3775:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface BBAK3775:fte-0/1/0.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: EE3093:fte-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface EE3093:fte-0/1/0.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: EE3093:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

  Logical interface EE3093:fte-0/1/2.32768
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    eth-switch

Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

```

Logical interface IC-WS001:fte-0/0/0.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/4.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/6.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/13.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
 Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
 Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/15.32768
 Flags: SNMP-Traps 0x0 Encapsulation: ENET2
 eth-switch

```

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/2.32768
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  eth-switch

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/7.32768
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  eth-switch

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/10.32768
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  eth-switch

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
Speed: Unspecified
  Device flags      : Present Running

Logical interface IC-WS001:bme0.0
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  inet  128.0.32.0      --> 0/0

Logical interface IC-WS001:bme0.1
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  inet  128.0.0.7/2
        128.0.130.2/2

```

Logical interface IC-WS001:bme0.2

Flags: Encapsulation: ENET2

inet 128.0.0.7/8
 128.0.130.2/8
 169.254.128.7/16
 169.254.193.1/16

Physical interface: IC-WS001:bme1, 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 IC-WS001:bme1.0

Flags: Encapsulation: ENET2

inet 128.0.0.1/2
 128.0.0.4/2
 128.0.0.16/2
 128.0.0.17/2
 128.0.0.24/2
 128.0.0.25/2
 128.0.0.26/2
 128.0.0.28/2
 128.0.0.29/2
 128.0.0.31/2

tnp 0x4

Physical interface: IC-WS001:qfabric, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
 Speed: Unspecified

Device flags : Present Running

Interface flags: SNMP-Traps

Logical interface IC-WS001:qfabric.0

Flags: SNMP-Traps Encapsulation: ENET2

inet
 mpls
 eth-switch

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
 Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Logical interface IC-WS001:vcp0.32769

Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcpl, Enabled, Physical link is Up

Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Logical interface IC-WS001:vcpl.32768

Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp2, Enabled, Physical link is Up

Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Logical interface IC-WS001:vcp2.32768

Flags: LinkAddress 0-0 Encapsulation: ENET2


```

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
  Device flags      : Present Running

  Logical interface IC-WS001:vcp3.32768
    Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
  Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
  Device flags      : Present Running

  Logical interface IC-WS001:vcp4.32768
    Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: NW-NG-0:bme0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
Speed: Unspecified
  Device flags      : Present Running

  Logical interface NW-NG-0:bme0.0
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.1/2
          128.0.0.5/2
          128.0.32.0/2
    tnp   0x5

  Logical interface NW-NG-0:bme0.1
    Flags: LinkAddress 0-0 Encapsulation: ENET2
    inet  128.0.0.9/2
          128.0.128.4/2

  Logical interface NW-NG-0:bme0.2
    Flags: Encapsulation: ENET2
    inet  128.0.0.9/8
          128.0.128.68/8
          169.254.128.9/16
          169.254.192.34/16

Physical interface: NW-NG-0:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified
  Device flags      : Present Running

```

Interface flags: SNMP-Traps

Logical interface NW-NG-0:qfabric.0

Flags: SNMP-Traps Encapsulation: ENET2

inet

mpls

eth-switch

Physical interface: NW-NG-0:vcp0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Logical interface NW-NG-0:vcp0.32769

Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: NW-NG-0:vcpl, Enabled, Physical link is Up

Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed: 1000mbps

Device flags : Present Running

Logical interface NW-NG-0:vcpl.32768

Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: P2659-C:bme0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified, Speed: Unspecified

Device flags : Present Running

Logical interface P2659-C:bme0.0

Flags: LinkAddress 0-0 Encapsulation: ENET2

inet 128.0.0.1/2

128.0.0.16/2

128.0.32.0/2

tnp 0x10

Logical interface P2659-C:bme0.1

Flags: LinkAddress 0-0 Encapsulation: ENET2

inet 128.0.0.8/2

128.0.130.4/2

Logical interface P2659-C:bme0.2

Flags: Encapsulation: ENET2

inet 128.0.0.8/8

```

128.0.130.4/8
169.254.128.8/16
169.254.193.2/16

```

```

Physical interface: P2659-C:qfabric, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified

```

```

  Device flags      : Present Running
  Interface flags: SNMP-Traps

```

```

Logical interface P2659-C:qfabric.0
  Flags: SNMP-Traps Encapsulation: ENET2
  inet
  mpls
  eth-switch

```

```

Physical interface: P2659-C:vcp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:
1000mbps

```

```

  Device flags      : Present Running

```

```

Logical interface P2659-C:vcp0.32769
  Flags: LinkAddress 0-0 Encapsulation: ENET2

```

```

Physical interface: P2659-C:vcpl, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps

```

```

  Device flags      : Present Running

```

```

Logical interface P2659-C:vcpl.32768
  Flags: LinkAddress 0-0 Encapsulation: ENET2

```

```

Physical interface: P2659-C:vcp2, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps

```

```

  Device flags      : Present Running

```

```

Logical interface P2659-C:vcp2.32768
  Flags: LinkAddress 0-0 Encapsulation: ENET2

```

```

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled

```

```

  Device flags      : Present Running

```

Interface flags: SNMP-Traps Internal: 0x0

Logical interface P2659-C:fte-0/1/2.32768

Flags: SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: ptor-0:bme0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
Speed: Unspecified

Device flags : Present Running

Logical interface ptor-0:bme0.0

Flags: LinkAddress 0-0 Encapsulation: ENET2

inet 128.0.0.1/2

128.0.0.17/2

128.0.32.0/2

tnp 0x11

Logical interface ptor-0:bme0.1

Flags: LinkAddress 0-0 Encapsulation: ENET2

inet 128.0.0.16/2

128.0.130.18/2

Logical interface ptor-0:bme0.2

Flags: Encapsulation: ENET2

inet 128.0.0.16/8

128.0.130.18/8

169.254.128.16/16

169.254.193.9/16

Physical interface: ptor-0:qfabric, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified

Device flags : Present Running

Interface flags: SNMP-Traps

Logical interface ptor-0:qfabric.0

Flags: SNMP-Traps Encapsulation: ENET2

inet

mpls

eth-switch

Physical interface: ptor-0:vc0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:

```

1000mbps
  Device flags      : Present Running

  Logical interface ptor-0:vcp0.32769
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: ptor-0:vcpl, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
  Device flags      : Present Running

  Logical interface ptor-0:vcpl.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: ptor-0:vcpl, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: 70, MTU: 1496, Clocking: Unspecified, Speed:
1000mbps
  Device flags      : Present Running

  Logical interface ptor-0:vcpl.32768
    Flags: LinkAddress 0-0 Encapsulation: ENET2

```

show interfaces fabric detail

```
user@switch> show interfaces fabric detail
```

```

Physical interface: BBAK0372:fte-0/1/0, Enabled, Physical link is Up
  Interface index: 49165, SNMP ifIndex: 1212678666, Generation: 140
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues        : 12 supported, 12 maximum usable queues
  Hold-times        : Up 0 ms, Down 0 ms
  Current address: 84:18:88:d1:fa:1f, Hardware address: 84:18:88:d1:fa:1f
  Last flapped      : 2012-11-09 21:36:41 UTC (4d 00:23 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes      :          14256654          0 bps
    Output bytes     :          9618986          0 bps
    Input packets    :          90511          0 pps
    Output packets   :          60101          0 pps
  IPv6 transit statistics:

```

```

    Input  bytes :           0
    Output bytes :           0
    Input  packets:           0
    Output packets:           0
Egress queues: 12 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort           0              0              0
  3 fcoe                  0              0              0
  4 no-loss               0              0              0
  7 network-cont          0              0              0
  8 mcast                 0              0              0
Active alarms   : None
Active defects  : None

Logical interface BBAK0372:fte-0/1/0.32768 (Index 71) (SNMP ifIndex 1212678667)
(Generation 136)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes :      12450372
    Output bytes :      11986557
    Input  packets:      90510
    Output packets:      62750
  Local statistics:
    Input  bytes :      12450372
    Output bytes :      11986557
    Input  packets:      90510
    Output packets:      62750
  Transit statistics:
    Input  bytes :           0          0 bps
    Output bytes :           0          0 bps
    Input  packets:           0          0 pps
    Output packets:           0          0 pps
  Protocol eth-switch, MTU: 0, Generation: 163, Route table: 0

```

show interfaces fabric extensive

user@switch> show interfaces fabric extensive

```

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Up
  Interface index: 49176, SNMP ifIndex: 1209008655, Generation: 155
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,

```

```

Flow control: Disabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
CoS queues    : 12 supported, 12 maximum usable queues
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:00:00:00:00:06, Hardware address: 00:00:00:00:00:06
Last flapped  : 2012-11-13 23:53:30 UTC (00:53:20 ago)
Statistics last cleared: Never

Traffic statistics:
  Input bytes   :                91179                0 bps
  Output bytes  :          361268221791          952985992 bps
  Input packets:                590                0 pps
  Output packets:          2580487185          850880 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: 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: 12 supported, 5 in use

Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 fabric_fcset      0                0                0
  1 fabric_fcset      0                0                0
  2 fabric_fcset      0                0                0
  3 fabric_fcset      0                0                0
  4 fabric_fcset      0                0                0
  5 fabric_fcset      0                0                0
  6 fabric_fcset      0                0                0
  7 fabric_fcset      0                0                0
  8 fabric_fcset      0          2582632925          0
  9 fabric_fcset      0                0                0
 10 fabric_fcset      0                0                0
 11 fabric_fcset      0                0                0

Active alarms   : None
Active defects  : None

MAC statistics:      Receive      Transmit

```

Total octets	91179	361268221791
Total packets	590	2580487185
Unicast packets	590	2580487185
Broadcast packets	0	0
Multicast packets	0	0
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	

MAC Priority Flow Control Statistics:

Priority : 0	0	0
Priority : 1	0	0
Priority : 2	0	0
Priority : 3	0	0
Priority : 4	0	0
Priority : 5	0	0
Priority : 6	0	0
Priority : 7	0	0

Packet Forwarding Engine configuration:

Destination slot: 0

Direction : Output

CoS transmit queue		Bandwidth		Buffer Priority	
Limit					
	%	bps	%	usec	
0 best-effort	5	2000000000	5	0	low
none					
3 fcoe	35	14000000000	35	0	low
none					
4 no-loss	35	14000000000	35	0	low
none					
7 network-control	5	2000000000	5	0	low
none					
8 mcast	20	8000000000	20	0	low
none					

Logical interface IC-WS001:fte-0/0/6.32768 (Index 85) (SNMP ifIndex 1209008656)
(Generation 150)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:


```

    Input  bytes  :                79496
    Output bytes  :            179860
    Input  packets:                590
    Output packets:                948
Local statistics:
    Input  bytes  :                79496
    Output bytes  :            179860
    Input  packets:                590
    Output packets:                948
Transit statistics:
    Input  bytes  :                0                0 bps
    Output bytes  :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
Protocol eth-switch, MTU: 0, Generation: 178, Route table: 0

```

show interfaces fabric terse

user@switch> **show interfaces fabric terse**

Item	Identifier	Connection	Configuration
Node group			
	BBAK3775	Connected	Configured
	NW-NG-0	Connected	Configured
	P2659-C	Connected	Configured
	ptor-0	Connected	Configured
Fabric control			
	FC-0	Connected	Configured
	FC-1	Connected	Configured

show interfaces fabric device-name

user@switch> **show interfaces fabric IC-WS001:fte-0/0/13**

```

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Up
  Interface index: 49177, SNMP ifIndex: 1209008767
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues        : 12 supported, 12 maximum usable queues

```

Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d
Last flapped : 2012-11-13 23:55:15 UTC (00:55:38 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
Active alarms : None
Active defects : None

Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1209008768)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Input packets : 748
Output packets: 954
Protocol eth-switch, MTU: 0

show interfaces ge

List of Syntax

[Syntax \(EX Series\) on page 1075](#)

[Syntax \(QFX Series\) on page 1075](#)

Syntax (EX Series)

```
show interfaces ge-fpc/pic/port
<brief | detail | extensive | terse>
<media>
<statistics>
```

Syntax (QFX Series)

```
show interfaces device-name:type-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display status information about the specified Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.

NOTE: You must have a transceiver plugged into an SFP or SFP+ port before information about the interface can be displayed.

Options

ge-fpc/pic/port—(EX Series) Display standard information about the specified Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) (EX Series) Display the specified level of output.

media—(Optional) (EX Series) Display media-specific information about network interfaces.

statistics—(Optional) (EX Series) Display static interface statistics.

brief | detail | extensive | terse—(Optional) (QFX Series) Display the specified level of output.

device-name:type-fpc/pic/port— (QFX Series) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.

descriptions—(Optional) (QFX Series) Display interface description strings.

media—(Optional) (QFX Series) Display media-specific information about network interfaces.

routing instance (all | instance-name)—(Optional) (QFX Series) Display the name of an individual routing-instance or display all routing-instances.

snmp-index snmp-index—(Optional) (QFX Series) Display information for the specified SNMP index of the interface.

statistics—(Optional) (QFX Series) Display static interface statistics.

Required Privilege Level

view

RELATED DOCUMENTATION

Monitoring Interface Status and Traffic

Troubleshooting Network Interfaces on EX3200 Switches

Troubleshooting Network Interfaces on EX4200 Switches

[Troubleshooting an Aggregated Ethernet Interface | 279](#)

[Junos OS Ethernet Interfaces Configuration Guide](#)

[Monitoring Interface Status and Traffic | 392](#)

[Troubleshooting Network Interfaces | 399](#)

Troubleshooting an Aggregated Ethernet Interface

Junos OS Network Interfaces Library for Routing Devices

List of Sample Output

[show interfaces ge-0/0/0 on page 1086](#)

[show interfaces ge-0/0/0 brief on page 1086](#)

[show interfaces ge-0/0/0 brief \(with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces\) on page 1087](#)

[show interfaces ge-0/0/0 detail on page 1087](#)

[show interfaces ge-0/0/4 extensive on page 1088](#)

Output Fields

Table 84 on page 1077 lists the output fields for the **show interfaces ge-** command. Output fields are listed in the approximate order in which they appear.

Table 84: show interfaces ge- Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface: Enabled or Disabled .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Description	Optional user-specified description.	brief detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. Default is 1514.	All levels
Speed	Speed of the interface: Auto if autonegotiation of speed is enabled; speed in megabits per second if the interface speed is explicitly configured.	All levels
Duplex	Link mode of the interface: Auto if autonegotiation of link mode is enabled; Full-Duplex or Half-Duplex if the link mode is explicitly configured.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	Autonegotiation status: Enabled or Disabled .	All levels

Table 84: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Remote-fault	Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
IEEE 802.3az Energy Efficient Ethernet	IEEE 802.3az Energy Efficient Ethernet status: Enabled or Disabled (appears only for EEE-capable Base-T copper Ethernet interfaces).	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	MAC address of the hardware.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled on the switch.</p>	detail extensive

Table 84: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 84: show interfaces ge- Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 84: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain time, it is promoted to an alarm. Based on the switch configuration, a defect can activate the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	<p>detail extensive none</p>

Table 84: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
Filter Statistics	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.	extensive

Table 84: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Complete—The autonegotiation process between the local and remote Ethernet interfaces was successful. • Incomplete—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation. • No autonegotiation—Local Ethernet interface has autonegotiation disabled and the link mode and speed are manually configured. • Link partner—Information from the link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. If the link mode of the remote device cannot be determined, the value is Unknown. • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, the types are: Symmetric (link partner supports PAUSE on receive and transmit); Asymmetric (link partner supports PAUSE on transmit); and Symmetric/Asymmetric (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Link partner speed—Speed of the link partner. • Local resolution—Resolution of the autonegotiation process on the local interface: <ul style="list-style-type: none"> • Flow control—Type of flow control that is used by the local interface. For Gigabit Ethernet interfaces, the types are: Symmetric (link partner supports PAUSE on receive and transmit); Asymmetric (link partner supports PAUSE on transmit); and Symmetric/Asymmetric (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). • Link mode—Link mode of local interface: either Full-duplex or Half-duplex. Displayed when Negotiation status is Incomplete. • Local link speed—Speed of the local interface. Displayed when Negotiation status is Incomplete. • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive

Table 84: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> • Destination slot—FPC slot number: <ul style="list-style-type: none"> • On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0. • On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch. • On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	detail extensive none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface. NOTE: For logical interfaces on EX Series switches, the traffic statistics fields in show interfaces commands show only control traffic; the traffic statistics do not include data traffic.	detail extensive
IPv6 transit statistics	EX Series switches do not support the collection and reporting of IPv6 transit statistics.	extensive
Local statistics	Number and rate of bytes and packets destined to and from the switch.	extensive

Table 84: show interfaces ge- Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive none
Input Filters	Names of any input filters applied to this interface.	detail extensive
Output Filters	Names of any output filters applied to this interface.	detail extensive
Flags	Information about protocol family flags. If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.	detail extensive
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flags.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces ge-0/0/0

user@switch> show interfaces ge-0/0/0

```
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 21
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
  Last flapped   : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 22)
  Flags: SNMP-Traps
  Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch
  Flags: None
```

show interfaces ge-0/0/0 brief

user@switch> show interfaces ge-0/0/0 brief

```
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Description: voice priority and tcp and icmp traffic rate-limiting filter at i
  ngress port
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
```

```
Link flags      : None
```

```
Logical interface ge-0/0/0.0
```

```
Flags: Device-Down SNMP-Traps Encapsulation: ENET2
eth-switch
```

show interfaces ge-0/0/0 brief (with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces)

```
user@switch> show interfaces ge-0/0/0 brief
```

```
Physical interface: ge-0/0/0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online,
IEEE 802.3az Energy Efficient Ethernet: Enabled, NO LPI
Device flags      : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags        : None
```

show interfaces ge-0/0/0 detail

```
user@switch> show interfaces ge-0/0/0 detail
```

```
Physical interface: ge-0/0/0, Enabled, Physical link is Up
Interface index: 193, SNMP ifIndex: 206, Generation: 196
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags        : None
CoS queues        : 8 supported, 8 maximum usable queues
Hold-times        : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped      : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes       :                               0          0 bps
Output bytes      :                               0          0 bps
Input packets     :                               0          0 pps
Output packets    :                               0          0 pps
IPv6 transit statistics:
```

```

Input  bytes :                0
Output bytes :                0
Input  packets:              0
Output packets:              0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets
  0 best-effort             0                0                0
  1 assured-forw            0                0                0
  5 expedited-fo            0                0                0
  7 network-cont            0                0                0
Active alarms   : None
Active defects  : None

Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
  Input  bytes :                0
  Output bytes :                0
  Input  packets:              0
  Output packets:              0
Local statistics:
  Input  bytes :                0
  Output bytes :                0
  Input  packets:              0
  Output packets:              0
Transit statistics:
  Input  bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input  packets:              0                0 pps
  Output packets:              0                0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,,

```

show interfaces ge-0/0/4 extensive

user@switch> show interfaces ge-0/0/4 extensive

```

Physical interface: ge-0/0/4, Enabled, Physical link is Up
Interface index: 165, SNMP ifIndex: 152, Generation: 168
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,

```



```

Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:33:65:44, Hardware address: 00:1f:12:33:65:44
Last flapped   : 2008-09-17 11:02:25 UTC (16:32:54 ago)
Statistics last cleared: Never

Traffic statistics:
  Input  bytes   :                0                0 bps
  Output bytes   :          2989761          984 bps
  Input  packets :                0                0 pps
  Output packets :          24307         1 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: 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 assured-forw     0                0                0
  5 expedited-fo     0                0                0
  7 network-cont     0                24307             0

Active alarms   : None
Active defects  : None

MAC statistics:
  Receive      Transmit
  Total octets  0          2989761
  Total packets 0          24307
  Unicast packets 0          0
  Broadcast packets 0          0
  Multicast packets 0          24307
  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
Code violations                  0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: None, Remote fault: OK,
    Link partner Speed: 1000 Mbps
  Local resolution:
    Flow control: None, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
  Direction : Output
  CoS transmit queue            Bandwidth            Buffer Priority
Limit                           %            bps      %            usec
0 best-effort                   95      950000000    95            NA      low
none
7 network-control               5       50000000     5            NA      low
none

Logical interface ge-0/0/4.0 (Index 82) (SNMP ifIndex 184) (Generation 147)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input  bytes :                0
  Output bytes :             4107883
  Input  packets:                0
  Output packets:             24307
IPv6 transit statistics:
  Input  bytes :                0
  Output bytes :                0
  Input  packets:                0
  Output packets:                0
Local statistics:
  Input  bytes :                0
  Output bytes :             4107883
  Input  packets:                0
  Output packets:             24307
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 eth-switch, Generation: 159, Route table: 0
  Flags: None
  Input Filters: f2,
  Output Filters: f1,,,
```

show interfaces (GRE)

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.

Command introduced in Junos OS Release 12.1 for EX Series switches.

Command introduced in Junos OS Release 13.2 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Command introduced before Junos OS Release 17.3R1.

Description

Display status information about the specified generic routing encapsulation (GRE) interface.

Options

interface-type—On M Series and T Series routers and EX Series switches, the interface type is ***gr-fpc/pic/port***.

brief | detail | extensive | terse—(Optional) Display the specified output level of interface information.

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.

NOTE: You can configure generic routing encapsulation (GRE) interfaces (*gre-x/y/z*) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the *MPLS Applications User Guide*.

Required Privilege Level

view

List of Sample Output

[show interfaces \(GRE\) on page 1097](#)

[show interfaces brief \(GRE\) on page 1097](#)

[show interfaces detail \(GRE\) on page 1098](#)

[show interfaces \(Layer 2 Services Over GRE Interfaces\) on page 1099](#)

[show interfaces extensive \(Layer 2 Services Over GRE Interfaces\) on page 1099](#)

[show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 1100](#)

[show interfaces extensive \(GRE\) on page 1102](#)

[show interfaces gr-2/0/10 for GRE IPv6 tunnel on page 1102](#)

Output Fields

[Table 85 on page 1093](#) lists the output fields for the **show interfaces** (GRE) command. Output fields are listed in the approximate order in which they appear.

Table 85: GRE show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 779 .	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Device Flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 779 .	All levels

Table 85: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Interface Flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 779 .	All levels
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support.	detail extensive
Flags	<p>Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under “Common Output Fields Description” on page 779. describe general information about the logical interface.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> • Reassemble-Pkts—If the Flags field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation. 	All levels

Table 85: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
IP-Header	<p>IP header of the logical interface. If the tunnel key statement is configured, this information is included in the IP Header entry.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> • df—If the IP-Header field includes this string immediately following the 16 bits of identification information (that is, if :df: displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation. 	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
L2 Routing Instance	Name of the Layer 2 routing instance associated with the GRE interface.	All levels
L3 Routing Instance	Name of the Layer 3 routing instance associated with the GRE interface.	All levels
Copy-tos-to-outer-ip-header	<p>Status of type of service (ToS) bits in the GRE packet header:</p> <ul style="list-style-type: none"> • On—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel. • Off—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header. <p>NOTE: EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always Off in switch output.</p>	detail extensive
Gre keepalives configured	<p>Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel.</p> <p>NOTE: EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always Off in switch output.</p>	detail extensive
Gre keepalives adjacency state	Status of the other end of the GRE tunnel: Up or Down . If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up.	detail extensive
Input packets	Number of packets received on the logical interface.	None specified

Table 85: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input rate—Rate of bits and packets received on the interface. • Output rate—Rate of bits and packets transmitted on the interface. 	detail extensive
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	detail extensive
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	detail extensive none
Protocol	Protocol family configured on the logical interface, such as iso , inet6 , or mpls .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none

Table 85: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (GRE)

user@host> show interfaces gr-1/2/0

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

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 192.0.2.2:192.0.2.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL

Input packets : 0
Output packets: 0
  Protocol inet, MTU: 1476
    Flags: None
    Addresses, Flags: Is-Primary
      Local: 198.51.100.1
```

show interfaces brief (GRE)

user@host> show interfaces gr-1/2/0 brief

```
Physical interface: gr-1/2/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface gr-1/2/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
  IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  inet 10.100.0.1/30
  mpls
```

show interfaces detail (GRE)

user@host> **show interfaces gr-1/2/0 detail**

```
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26, Generation: 13
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags     : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 192.0.2.2:192.0.2.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL

  Traffic statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Local statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Transit statistics:
```

```

Input  bytes   :                0                0 bps
Output bytes   :                0                0 bps
Input  packets :                0                0 pps
Output packets :                0                0 pps
Protocol inet, MTU: 1476, Generation: 12, Route table: 0
  Flags: None
  Addresses, Flags: Is-Primary
    Destination: Unspecified, Local: 198.51.100.1, Broadcast: Unspecified,
    Generation: 15

```

show interfaces (Layer 2 Services Over GRE Interfaces)

user@host> **show interfaces gr-2/2/10**

```

show interfaces gr-2/2/10
Physical interface: gr-2/2/10, Enabled, Physical link is Up
  Interface index: 214, SNMP ifIndex: 690
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate        : 0 bps (0 pps)
  Output rate       : 0 bps (0 pps)

  Logical interface gr-2/2/10.0 (Index 342) (SNMP ifIndex 10834)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
  203.0.113.1:203.0.113.254:47:df:64:0000000000000000 Encapsulation: GRE-NULL
    L2 Routing Instance: vs1, L3 Routing Instance: default
    Copy-tos-to-outer-ip-header: Off
    Gre keepalives configured: Off, Gre keepalives adjacency state: down
    Input packets : 2
    Output packets: 0
    Protocol bridge, MTU: 1476
      Flags: Sendbroadcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 6/8, Local: 6.0.0.1, Broadcast: 6.255.255.255

```

show interfaces extensive (Layer 2 Services Over GRE Interfaces)

user@host> **show interfaces gr-2/2/10.0 extensive**

```

Flags: SNMP-Traps Encapsulation: ENET2

```

```

L2 Routing Instance: vs1, L3 Routing Instance: default
Traffic statistics:
  Input  bytes   :          58851250
  Output bytes   :              0
  Input  packets :        1279375
  Output packets :              0
Local statistics:
  Input  bytes   :              0
  Output bytes   :              0
  Input  packets :              0
  Output packets :              0
Transit statistics:
  Input  bytes   :          58851250          75136 bps
  Output bytes   :              0              0 bps
  Input  packets :        1279375          204 pps
  Output packets :              0              0 pps
Protocol bridge, MTU: 1476, Generation: 175, Route table: 7
Flags: Access-Mode

```

show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch

user@host> show interfaces gr-2/0/15 detail

```

Physical interface: gr-2/0/15, Enabled, Physical link is Up
  Interface index: 195, SNMP ifIndex: 846, Generation: 198
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:d2, Hardware address: 00:00:5e:00:53:d2
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: 2011-09-14 17:43:15 UTC (00:00:18 ago)
Traffic statistics:
  Input  bytes   :          5600636          0 bps
  Output bytes   :          5600636          0 bps
  Input  packets :         20007          0 pps
  Output packets :         20007          0 pps
IPv6 transit statistics:
  Input  bytes   :              0
  Output bytes   :              0
  Input  packets :              0
  Output packets :              0

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)

```

```

Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 192.168.30.2:192.168.20.3:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
  Input  bytes   :           5600886
  Output bytes   :           2881784
  Input  packets :           20010
  Output packets :           10018
Local statistics:
  Input  bytes   :             398
  Output bytes   :             264
  Input  packets :              5
  Output packets :              3
Transit statistics:
  Input  bytes   :           5600488           0 bps
  Output bytes   :           2881520           0 bps
  Input  packets :           20005           0 pps
  Output packets :           10015           0 pps
Protocol inet, Generation: 159, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.10.10/8, Local: 10.10.10.10, Broadcast: 10.10.10.255,
    Generation: 144

Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 192.168.40.2:192.168.30.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
  Input  bytes   :             260
  Output bytes   :          2880148
  Input  packets :              4
  Output packets :          10002
Local statistics:
  Input  bytes   :             112
  Output bytes   :              0
  Input  packets :              2
  Output packets :              0
Transit statistics:

```

```

Input  bytes   :                148                0 bps
Output bytes   :            2880148                0 bps
Input  packets :                2                0 pps
Output packets :            10002                0 pps
Protocol inet, Generation: 171, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.10.10/8, Local: 10.10.10.10, Broadcast: 10.10.10.255,
  Generation: 160

```

show interfaces extensive (GRE)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 1098](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 1100](#).

show interfaces gr-2/0/10 for GRE IPv6 tunnel

user@host> **show interfaces gr-2/0/10**

```

show interfaces gr-2/0/10
Physical interface: gr-2/0/10, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 559
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 4952 bps (3 pps)
  Output rate    : 200 bps (0 pps)

Logical interface gr-2/0/10.0 (Index 355) (SNMP ifIndex 857)
  Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
1000::11:0:11:1-1000::11:2:13:2-47-64-0-0-0000000000000000 Encapsulation: GRE-NULL

Copy-tos-to-outer-ip-header: Off, Copy-tos-to-outer-ip-header-transit: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Input packets : 60
Output packets: 83
Protocol inet, MTU: 9082
Max nh cache: 0, New hold nh limit: 0, 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: 14.0.13/24, Local: 14.0.13.1, Broadcast: 14.0.13.255
  Protocol iso, MTU: 9082

```

```
Protocol inet6, MTU: 9082
Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
NH drop cnt: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 1400::14:0:13:0/120, Local: 1400::14:0:13:1
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::2a0:a520:2875:4992
Protocol mpls, MTU: 9070, Maximum labels: 3
  Flags: Is-Primary
```

show interfaces irb

Syntax

```
show interfaces irb
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance instance-name>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced in Junos OS Release 12.3R2.

Command introduced in Junos OS Release 12.3R2 for EX Series switches.

Command introduced in Junos OS Release 13.2 for the QFX Series

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.

media—(Optional) Display media-specific information about network interfaces.

routing-instance *instance-name*—(Optional) Display information for the interface with the specified SNMP index.

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 VLAN that has a Layer 3 protocol configured.

Required Privilege Level

view

List of Sample Output

[show interfaces irb extensive on page 1110](#)

[show interfaces irb snmp-index on page 1111](#)

Output Fields

Table 86 on page 1105 lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

Table 86: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 779 .	All levels
Proto	Protocol configured on the interface.	terse
Interface index	Physical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Type	Physical interface type.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	detail extensive brief none
MTU	MTU size on the physical interface.	detail extensive brief none
Clocking	Reference clock source: Internal or External . Always unspecified on IRB interfaces.	detail extensive brief
Speed	Speed at which the interface is running. Always unspecified on IRB interfaces.	detail extensive brief
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 779 .	detail extensive brief none
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 779 .	detail extensive brief none

Table 86: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link type	Physical interface link type: full duplex or half duplex .	detail extensive none
Link flags	Information about the link. Possible values are described in the “Links Flags” section under “Common Output Fields Description” on page 779 .	detail extensive none
Physical Info	Physical interface information.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	MAC address of the hardware.	detail extensive none
Alternate link address	Backup address of the link.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 86: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	detail extensive

Table 86: show interfaces irb Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface (which reflects its initialization sequence).	detail extensive none
SNMP ifIndex	SNMP interface index number of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 779 .	detail extensive
Encapsulation	Encapsulation on the logical interface.	detail extensive
Bandwidth	Dummy value that is ignored by an IRB interface. IRB interfaces are pseudo interfaces and do not have physical bandwidth associated with them.	detail extensive
Routing Instance	Routing instance IRB is configured under.	detail extensive
Bridging Domain	Bridging domain IRB is participating in.	detail extensive

Table 86: show interfaces irb Output Fields (continued)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the logical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. 	detail extensive
IPv6 transit statistics	<p>Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine.	detail extensive
Transit statistics	Statistics for traffic transiting the router.	detail extensive
Protocol	Protocol family configured on the local interface. Possible values are described in the “Protocol Field” section under “Common Output Fields Description” on page 779 .	detail extensive
MTU	Maximum transmission unit size on the logical interface.	detail extensive
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive
Addresses, Flags	Information about address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 779 .	detail extensive
Policer	The policer that is to be evaluated when packets are received or transmitted on the interface.	detail extensive

Table 86: show interfaces irb Output Fields (continued)

Field Name	Field Description	Level of Output
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under “Common Output Fields Description” on page 779.	detail extensive

Sample Output

show interfaces irb extensive

user@host> **show interfaces irb extensive**

```
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: Unspecified
  Device flags      : Present Running
  Interface flags: SNMP-Traps
  Link type        : Full-Duplex
  Link flags       : None
  Physical info    : Unspecified
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped     : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  IPv6 transit statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
```

0

```

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Traffic statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  IPv6 transit statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Local statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Transit statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
  IPv6 transit statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Protocol inet, MTU: 1500, Generation: 154, Route table: 0
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.1/8, Local: 10.51.1.2, Broadcast: 10.51.1.255,
      Generation: 155
  Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
    Flags: Is-Primary
    Policers: Input: __default_arp_policer

```

show interfaces irb snmp-index

user@host> **show interfaces irb snmp-index 25**

```

Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags      : Present Running
  Interface flags: SNMP-Traps
  Link type         : Full-Duplex
  Link flags        : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped     : Never
    Input packets : 0
    Output packets: 0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1500
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.1/8, Local: 10.51.1.2, Broadcast: 10.51.1.255
  Protocol multiservice, MTU: 1500
    Flags: Is-Primary

```


show interfaces mc-ae

Syntax

```
show interfaces mc-ae id identifier unit number
```

Release Information

Command introduced in Junos OS Release 9.6 for the MX Series.

Command introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Configuration Consistency Check output field added in Junos OS Release 15.1X53-D60 for the QFX Series.

Description

On peers with multichassis aggregated Ethernet (**mc-aeX**) interfaces, use this command to display information about the multichassis aggregated Ethernet interfaces.

NOTE: In Junos OS Release 17.4R1, this command is not supported on EX4300, EX9200, PTX10000, QFX10002, and QFX10008 devices.

Options

id *identifier*—(Optional) Specify the name of the multichassis aggregated Ethernet interface.

unit *number*—(Optional) Specify the logical interface by unit number.

Required Privilege Level

view

List of Sample Output

[show interfaces mc-ae \(EX Series \) on page 1115](#)

[show interfaces mc-ae \(MX Series\) on page 1115](#)

[show interfaces mc-ae \(Active/Active Bridging and VRRP over IRB on MX Series\) on page 1115](#)

Output Fields

[Table 87 on page 1114](#) lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

Table 87: show interfaces mc-ae Output Fields

Output Field Name	Field Description
Current State Machine's State	Specifies the state of the MC-LAG initialization state machine.
Configuration Consistency Check	Specifies the status of the MC-LAG configuration consistency check feature. The status is either Passed or Failed . If the status is Failed , the system will display the name of the parameter that failed consistency check. If there are multiple inconsistencies, only the first inconsistency is shown. If the enforcement level for the MC-LAG parameter was mandatory, and you did not configure that parameter correctly, the command will show that the MC-LAG interface is down.
Member Link	Specifies the identifiers of the configured multichassis link aggregated interface members.
Local Status	Specifies the status of the local link: active or standby .
Peer Status	Specifies the status of the peer link: active or standby .
Peer State	Specifies the status of the local and peer links in an active/active MC-LAG configuration.
Logical Interface	Specifies the identifier and unit of the AE interface.
Topology Type	Specifies the bridge configured on the AE.
Local State	Specifies if the local device is up or down.
Peer State	Specifies if the peer device is up or down.
Peer Ip/MCP/State	Specifies the multichassis protection (MCP) link or the interchassis link-protection link (ICL-PL) for all of the multichassis aggregated Ethernet interfaces that are part of the peer.

Sample Output

show interfaces mc-ae (EX Series)

user@switch> show interfaces mc-ae ae1 512

```
Member Link           : ae1
Current State Machine's State: mcae active state
Configuration Consistency Check : Failed (redundancy group id mismatch)
Local Status          : active
Local State           : up
Peer Status           : standby
Peer State            : up
    Logical Interface   : ae1.0
    Topology Type       : bridge
    Local State         : up
    Peer State          : up
    Peer Ip/MCP/State   : 10.1.1.1 ae0.0 up
```

show interfaces mc-ae (MX Series)

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)

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 me0

Syntax

```
show interfaces me0  
<brief | detail | extensive | terse>  
<descriptions>  
<media>  
<routing-instance>  
<statistics>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display status information about the management Ethernet interface.

Options

none—Display standard information about the management 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.

routing-instance—(Optional) Display the name of the routing instance.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

RELATED DOCUMENTATION

Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch

Configuring Firewall Filters (CLI Procedure)

List of Sample Output

[show interfaces me0 on page 1122](#)

[show interfaces me0 brief on page 1123](#)

[show interfaces me0 detail on page 1123](#)

[show interfaces me0 extensive on page 1125](#)

Output Fields

Table 88 on page 1118 lists the output fields for the **show interfaces me0** command. Output fields are listed in the approximate order in which they appear.

Table 88: show interfaces me0 Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface: Enabled or Disabled .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Description	Optional user-specified description.	brief detail extensive
Type	Information about the type of functional interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. The default is 1514.	All levels
Clocking	Interface that acts as a clock source. This field is not supported on EX Series switches and the default value is always Unspecified .	detail extensive
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link type	Information about whether the link is duplex and whether the negotiation is manual or automatic.	detail extensive none

Table 88: show interfaces me0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
Physical info	Information about the device dependent physical interface selector. This field is applied only when a clocking option is specified. This field is not supported on EX Series switches and the default value is always Unspecified .	detail extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	MAC address of the hardware.	detail extensive none
Alternate link address	Information about alternate hardware address.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour:minute:second timezone (weeksw:daysdhour:minute:second ago)</i> . For example, Last flapped: 2008-01-16 10:52:40 UTC (3w:3d 22:58 ago) .	detail extensive none
Statistics last cleared	Time when the statistics for the interface was last set to zero. The format is Last flapped: <i>year-month-day hour:minute:second timezone (weeksw:daysdhour:minute:second ago)</i> . For example, Last flapped: 2008-01-16 10:52:40 UTC (3w:3d 22:58 ago) .	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <p>Following are fields in Traffic statistics:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 88: show interfaces me0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>Number and rate of bytes and IPv6 packets received and transmitted on the physical interface.</p> <p>Following are fields in IPv6 transit statistics:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and frame checksum (FCS) errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid FCS. • Runts—Number of frames received that are smaller than the runt threshold. • Giants— Number of packets that exceed the size for the medium. For example, if the medium is Ethernet, the Giant field shows the count of packets with size greater than 1518 bytes. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • Resource errors—Sum of transmit drops. 	extensive

Table 88: show interfaces me0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly. It increases 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 increment 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. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	detail extensive
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	detail extensive
Local statistics	Number and rate of bytes and packets destined to and exiting from the switch.	extensive

Table 88: show interfaces me0 Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protocol	Protocol family.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive
Flags	Information about protocol family flags.	detail extensive
Input Filter	Ingress filter name.	extensive
Output Filter	Egress filter name.	extensive
Addresses	Information about the management interface addresses.	detail extensive none
Flags	Information about the address flags.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces me0

```
user@switch> show interfaces me0
```

```
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
```

```

Interface flags: SNMP-Traps
Link type      : Full-Duplex
Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:32 ago)
  Input packets : 1661830
  Output packets: 3200

Logical interface me0.0 (Index 3) (SNMP ifIndex 34)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 1661830
  Output packets: 3200
  Protocol inet
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.204.32/20, Local: 10.204.33.103,
      Broadcast: 10.204.47.255
  Protocol inet6
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf

```

show interfaces me0 brief

user@switch> **show interfaces me0 brief**

```

Physical interface: me0, 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 me0.0
  Flags: SNMP-Traps Encapsulation: ENET2
  inet  10.204.33.103/20
  inet6 fe80::21f:12ff:fe35:3cbf/64

```

show interfaces me0 detail

user@switch> **show interfaces me0 detail**

```

Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, Generation: 1

```

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: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf

Alternate link address: Unspecified

Last flapped : 2010-07-31 23:45:50 PDT (5d 00:37 ago)

Statistics last cleared: Never

Traffic statistics:

Input bytes : 366663167

Output bytes : 498590

Input packets: 1664031

Output packets: 3259

IPv6 transit statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

Input bytes : 366665637

Output bytes : 500569

Input packets: 1664048

Output packets: 3275

IPv6 transit statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 366665637

Output bytes : 500569

Input packets: 1664048

Output packets: 3275

Protocol inet, Generation: 1, Route table: 0

Flags: Is-Primary

Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.204.32/20, Local: 10.204.33.103, Broadcast: 10.204.47.255,

Generation: 1

```

Protocol inet6, Generation: 2, Route table: 0
  Flags: Is-Primary
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf
  Generation: 2

```

show interfaces me0 extensive

user@switch> show interfaces me0 extensive

```

Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, 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: 00:1f:12:38:58:bf, Hardware address: 00:1f:12:38:58:bf
  Alternate link address: Unspecified
  Last flapped   : 2010-08-15 06:27:33 UTC (03:06:22 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes   :           82310392
    Output bytes   :           1966952
    Input  packets :           110453
    Output packets :            17747
  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: 1, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:

```

```
Input  bytes :      82310392
Output bytes :      1966952
Input  packets:      110453
Output packets:      17747
Local statistics:
Input  bytes :      82310392
Output bytes :      1966952
Input  packets:      110453
Output packets:      17747
Protocol inet, Generation: 1, Route table: 0
  Flags: Is-Primary
  Input Filters: mgmt_filter,
  Addresses, Flags: Is-Default Is-Preferred Is-Primary
    Destination: 10.204.96/20, Local: 10.204.96.234,
    Broadcast: 10.204.111.255, Generation: 1
```

show interfaces queue

Syntax

```
show interfaces queue
<aggregate | remaining-traffic>
<both-ingress-egress>
<egress>
<forwarding-class forwarding-class>
<ingress>
< interface-name>
<l2-statistics>
<buffer-occupancy>
```

Release Information

Command introduced before Junos OS Release 7.4.

both-ingress-egress, **egress**, and **ingress** options introduced in Junos OS Release 7.6.

Command introduced in Junos OS Release 9.0 for EX Series.

Command introduced in Junos OS Release 11.1 for the QFX Series.

l2-statistics option introduced in Junos OS Release 12.1.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

buffer-occupancy statement introduced in Junos OS Release 19.1R1 for QFX5000 Series switches.

Description

Display class-of-service (CoS) queue information for physical interfaces.

Options

none—Show detailed CoS queue statistics for all physical interfaces.

aggregate—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

both-ingress-egress—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

egress—(Optional) Display egress queue statistics.

forwarding-class forwarding-class—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

ingress—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

interface-name—(Optional) Show detailed CoS queue statistics for the specified interface.

l2-statistics—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

buffer-occupancy—Displays the peak buffer occupancy for each queue while **buffer-monitor-enable** is enabled at the `[edit chassis fpc slot-number traffic-manager]` hierarchy level.

remaining-traffic—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.

Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 89 on page 1128](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 89 on page 1128](#).

Table 89: Layer 2 Overhead and Transmitted Packets or Byte Counts

Protocol	Fragmentation		LFI
	First fragmentation	Second to <i>n</i> fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

Layer 2 Statistics—Fragmentation Overhead Calculation

MLPPP/MC-MLPPP Overhead details:

=====

Fragment 1:

Outer PPP header : 4 bytes

Long or short sequence MLPPP header	: 4 bytes or 2 bytes
Inner PPP header	: 1 byte
HDLC flag and FCS bytes	: 4 bytes

Fragments 2 .. n :

Outer PPP header	: 4 bytes
Long or short sequence MLPPP header	: 4 bytes or 2 bytes
HDLC flag and FCS bytes	: 4 bytes

MLFR (FRF15) Overhead details:

=====

Fragment 1:

Framerelay header	: 2 bytes
Control,NLPID	: 2 bytes
Fragmentaion header	: 2 bytes
Inner proto	: 2 bytes
HDLC flag and FCS	: 4 bytes

Fragments 2 ...n :

Framerelay header	: 2 bytes
Control,NLPID	: 2 bytes
Fragmentaion header	: 2 bytes
HDLC flag and FCS	: 4 bytes

MFR (FRF16) Overhead details:

=====

Fragment 1:

Fragmentaion header	: 2 bytes
Framerelay header	: 2 bytes
Inner proto	: 2 bytes
HDLC flag and FCS	: 4 bytes

Fragments 2 ...n :

Fragmentaion header	: 2 bytes
Framerelay header	: 2 bytes
HDLC flag and FCS	: 4 bytes

Overhead with LFI

```
MLPPP(Long & short sequence):
=====
    Outer PPP header      : 4 bytes
    HDLC flag and FCS     : 4 bytes

MLFR (FRF15):
=====
    Framerelay header     : 2 bytes
    Control,NLPID         : 2 bytes
    HDLC flag and FCS     : 4 bytes
```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

remaining-traffic—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.

NOTE: For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For related CoS operational mode commands, see the [CLI Explorer](#).

Required Privilege Level

view

Release History Table

Release	Description
18.3R1	Starting with Junos OS 18.3R1, the Tail-dropped packets counter is supported on PTX Series Packet Transport Routers.
16.1	Starting with Junos OS Release 16.1, Last-packet enqueued output field is introduced.

RELATED DOCUMENTATION

Monitoring Interface Status and Traffic
Monitoring Interfaces That Have CoS Components
Defining CoS Schedulers and Scheduler Maps (CLI Procedure)

Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)

ATM Interfaces User Guide for Routing Devices

Layer 2 Bridging, Address Learning, and Forwarding User Guide

List of Sample Output

[show interfaces queue \(Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC\) on page 1140](#)

[show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 1141](#)

[show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 1143](#)

[show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 1144](#)

[show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 1149](#)

[show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 1153](#)

[show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 1156](#)

[show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 1158](#)

[show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 1160](#)

[show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 1163](#)

[show interfaces queue \(QFX Series\) on page 1170](#)

[show interfaces queue l2-statistics \(lsq interface\) on page 1171](#)

[show interfaces queue lsq \(lsq-ifd\) on page 1173](#)

[show interfaces queue \(Aggregated Ethernet on a MX series Router\) on page 1175](#)

[show interfaces queue ge-0/0/0 \(EX2200 Switch\) on page 1177](#)

[show interfaces queue xe-6/0/39 \(Line Card with Oversubscribed Ports in an EX8200 Switch\) on page 1178](#)

[show interfaces queue xe-0/0/2 buffer-occupancy \(QFX5000 Series switch\) on page 1181](#)

Output Fields

[Table 90 on page 1132](#) lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 90: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 779 .
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.

Table 90: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Ingress queues supported	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
Ingress queues in use	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
Output queues supported	Total number of output queues supported on the specified interface.
Output queues in use	Total number of output queues in use on the specified interface.
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. <p>NOTE: This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p> <ul style="list-style-type: none"> • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
Burst size	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.

The following output fields are applicable to both interface component and Packet Forwarding component in the **show interfaces queue** command:

Queue	Queue number.
Forwarding classes	Forwarding class name.

Table 90: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Queued Packets	<p>Number of packets queued to this queue.</p> <p>NOTE: For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see “Additional Information” on page 1130.</p> <p>NOTE: This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p> <p>This field is not supported on EX Series switches due to hardware limitations.</p>
Queued Bytes	<p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see Table 91 on page 1137.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see “Additional Information” on page 1130.</p> <p>NOTE: This field is not supported on QFX5100, QFX5110, QFX5200, and QFX5210 switches due to hardware limitations.</p> <p>This field is not supported on EX Series switches due to hardware limitations.</p>
Transmitted Packets	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.</p> <p>NOTE: For Layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 1128</p>
Transmitted Bytes	<p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see Table 91 on page 1137.</p> <p>NOTE: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p>NOTE: For Layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 1128</p>

Table 90: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Tail-dropped packets	<p>Number of packets dropped because of tail drop.</p> <p>NOTE: Starting with Junos OS 18.3R1, the Tail-dropped packets counter is supported on PTX Series Packet Transport Routers.</p>
RL-dropped packets	<p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see “Additional Information” on page 1130.</p> <p>NOTE: The RL-dropped packets counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.</p>
RL-dropped bytes	<p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see “Additional Information” on page 1130.</p>
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. • Low, TCP—Number of low-loss priority TCP packets dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. • High, TCP—Number of high-loss priority TCP packets dropped because of RED. • (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> • Low—Number of low-loss priority packets dropped because of RED. • Medium-low—Number of medium-low loss priority packets dropped because of RED. • Medium-high—Number of medium-high loss priority packets dropped because of RED. • High—Number of high-loss priority packets dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Table 90: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see Table 91 on page 1137.</p> <ul style="list-style-type: none"> • (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority TCP bytes dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
Queue-depth bytes	Displays queue-depth average, current, peak, and maximum values for RTP queues. Because queue-depth values cannot be aggregated, displays the values for RTP queues regardless of whether aggregate , remaining-traffic , or neither option is selected.
Peak	(QFX5000 Series switches only) Displays the peak buffer occupancy for the queue while buffer-monitor-enable is enabled at the <code>[edit chassis fpc slot-number traffic-manager]</code> hierarchy level.
Last-packet enqueued	Starting with Junos OS Release 16.1, Last-packet enqueued output field is introduced. If packet-timestamp is enabled for an FPC, shows the day, date, time, and year in the format <i>day-of-the-week month day-date hh:mm:ss yyyy</i> when a packet was enqueued in the CoS queue. When the timestamp is aggregated across all active Packet Forwarding Engines, the latest timestamp for each CoS queue is reported.

Byte counts vary by interface hardware. [Table 91 on page 1137](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 91 on page 1137](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 91: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	-

Table 91: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap. <p>M Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> • Queued: The sum of the transmitted bytes and the RED dropped bytes. • Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble. • RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble (does not include the VLAN header or MPLS pushed bytes). 	The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.

Table 91: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. <p>M Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes • RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet 	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	

Table 91: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	-
8-port 1G IQ2 and IQ2-E PICs		Transmitted: 478 bytes of Layer 3 packet.	

Sample Output

show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
```

```
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          :          131300649          141751 pps
    Bytes            :       11287964840       99793248 bps
  Transmitted:
```

```

Packets          :          131300649          141751 pps
Bytes            :          11287964840        99793248 bps
Tail-dropped packets :              0              0 pps
RL-dropped packets :          205050862        602295 pps
RL-dropped bytes  :          13595326612      327648832 bps
RED-dropped packets :              0              0 pps
  Low             :              0              0 pps
  Medium-low      :              0              0 pps
  Medium-high     :              0              0 pps
  High            :              0              0 pps
RED-dropped bytes :              0              0 bps
  Low             :              0              0 bps
  Medium-low      :              0              0 bps
  Medium-high     :              0              0 bps
  High            :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :              0              0 pps
    Bytes            :              0              0 bps

```

show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
```

```

Physical interface: ae1, Enabled, Physical link is Up
Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :              5              0 pps
    Bytes            :             242              0 bps
  Transmitted:
    Packets          :              5              0 pps
    Bytes            :             242              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets :              0              0 pps
    RED-dropped bytes  :              0              0 bps
Queue: 1, Forwarding classes: afl
  Queued:
    Packets          :          42603765        595484 pps

```

```

    Bytes          :          5453281920          609776496 bps
Transmitted:
    Packets        :          42603765          595484 pps
    Bytes          :          5453281920          609776496 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: efl
Queued:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps
Transmitted:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: nc
Queued:
    Packets        :          45          0 pps
    Bytes          :          3930          0 bps
Transmitted:
    Packets        :          45          0 pps
    Bytes          :          3930          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 4, Forwarding classes: afl1
Queued:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps
Transmitted:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 5, Forwarding classes: efl1
Queued:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps
Transmitted:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps

```

```

Tail-dropped packets :                0                0 pps
RED-dropped packets  :                0                0 pps
RED-dropped bytes   :                0                0 bps
Queue: 6, Forwarding classes: af12
  Queued:
    Packets          :                31296413            437436 pps
    Bytes            :                4005940864          447935200 bps
  Transmitted:
    Packets          :                31296413            437436 pps
    Bytes            :                4005940864          447935200 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps
Queue: 7, Forwarding classes: nc2
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps

```

show interfaces queue (Gigabit Ethernet on a T640 Router)

user@host> **show interfaces queue**

```

Physical interface: ge-7/0/1, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 42
  Forwarding classes: 8 supported, 8 in use
  Output queues: 8 supported, 8 in use
  Queue: 0, Forwarding classes: be
    Queued:
      Packets          :                13                0 pps
      Bytes            :                622                0 bps
    Transmitted:
      Packets          :                13                0 pps
      Bytes            :                622                0 bps
      Tail-dropped packets :                0                0 pps
      RED-dropped packets :                0                0 pps
      RED-dropped bytes   :                0                0 bps
  Queue: 1, Forwarding classes: af1

```

```

Queued:
  Packets      :      1725947945      372178 pps
  Bytes       :      220921336960    381110432 bps
Transmitted:
  Packets      :      1725947945      372178 pps
  Bytes       :      220921336960    381110432 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: ef1
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: nc
Queued:
  Packets      :      571      0 pps
  Bytes       :      49318     336 bps
Transmitted:
  Packets      :      571      0 pps
  Bytes       :      49318     336 bps
  Tail-dropped packets :      0      0 pps
  RED-dropped packets :      0      0 pps
  RED-dropped bytes  :      0      0 bps

```

show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

user@host> show interfaces queue ge-2/2/9 aggregate

```

Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets      :      148450735      947295 pps
  Bytes       :      8016344944    409228848 bps
Transmitted:

```



```

Packets          :          76397439          487512 pps
Bytes            :          4125461868        210602376 bps
Tail-dropped packets : Not Available
RED-dropped packets :          72053285          459783 pps
  Low            :          72053285          459783 pps
  Medium-low     :              0              0 pps
  Medium-high    :              0              0 pps
  High           :              0              0 pps
RED-dropped bytes :          3890877444        198626472 bps
  Low            :          3890877444        198626472 bps
  Medium-low     :              0              0 bps
  Medium-high    :              0              0 bps
  High           :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :              0              0 pps
    Bytes            :              0              0 bps
  Transmitted:
    Packets          :              0              0 pps
    Bytes            :              0              0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :              0              0 pps
      Low            :              0              0 pps
      Medium-low     :              0              0 pps
      Medium-high    :              0              0 pps
      High           :              0              0 pps
    RED-dropped bytes :              0              0 bps
      Low            :              0              0 bps
      Medium-low     :              0              0 bps
      Medium-high    :              0              0 bps
      High           :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets          :          410278257          473940 pps
    Bytes            :          22156199518        204742296 bps
  Transmitted:
    Packets          :          4850003           4033 pps
    Bytes            :          261900162        1742256 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :          405425693          469907 pps
      Low            :          405425693          469907 pps
      Medium-low     :              0              0 pps
      Medium-high    :              0              0 pps
      High           :              0              0 pps

```

```

RED-dropped bytes      :          21892988124          203000040 bps
  Low                  :          21892988124          203000040 bps
  Medium-low           :                   0              0 bps
  Medium-high          :                   0              0 bps
  High                 :                   0              0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets            :                   0              0 pps
    Bytes              :                   0              0 bps
  Transmitted:
    Packets            :                   0              0 pps
    Bytes              :                   0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets  :                   0              0 pps
    Low                :                   0              0 pps
    Medium-low         :                   0              0 pps
    Medium-high        :                   0              0 pps
    High               :                   0              0 pps
  RED-dropped bytes    :                   0              0 bps
    Low                :                   0              0 bps
    Medium-low         :                   0              0 bps
    Medium-high        :                   0              0 bps
    High               :                   0              0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets            :          76605230          485376 pps
    Bytes              :          5209211400        264044560 bps
  Transmitted:
    Packets            :          76444631          484336 pps
    Bytes              :          5198235612        263478800 bps
  Tail-dropped packets : Not Available
  RED-dropped packets  :          160475          1040 pps
    Low                :          160475          1040 pps
    Medium-low         :                   0              0 pps
    Medium-high        :                   0              0 pps
    High               :                   0              0 pps
  RED-dropped bytes    :          10912300          565760 bps
    Low                :          10912300          565760 bps
    Medium-low         :                   0              0 bps
    Medium-high        :                   0              0 bps
    High               :                   0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding

```

```

Queued:
  Packets      :                0                0 pps
  Bytes       :                0                0 bps
Transmitted:
  Packets      :                0                0 pps
  Bytes       :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
    Low       :                0                0 pps
    Medium-low :                0                0 pps
    Medium-high :                0                0 pps
    High      :                0                0 pps
  RED-dropped bytes :                0                0 bps
    Low       :                0                0 bps
    Medium-low :                0                0 bps
    Medium-high :                0                0 bps
    High      :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          4836136          3912 pps
  Bytes       :      333402032      2139056 bps
Transmitted:
  Packets      :          3600866          1459 pps
  Bytes       :      244858888      793696 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :          1225034          2450 pps
    Low       :          1225034          2450 pps
    Medium-low :                0                0 pps
    Medium-high :                0                0 pps
    High      :                0                0 pps
  RED-dropped bytes :          83302312      1333072 bps
    Low       :          83302312      1333072 bps
    Medium-low :                0                0 bps
    Medium-high :                0                0 bps
    High      :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :                0                0 pps
  Bytes       :                0                0 bps
Transmitted:
  Packets      :                0                0 pps
  Bytes       :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps

```

Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	77059796	486384 pps
Bytes	:	3544750624	178989576 bps

Transmitted:

Packets	:	77059797	486381 pps
Bytes	:	3544750670	178988248 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps

```

    Low                :                0                0 bps
    Medium-low         :                0                0 bps
    Medium-high        :                0                0 bps
    High               :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets           :            4846580            3934 pps
    Bytes              :        222942680        1447768 bps
  Transmitted:
    Packets           :            4846580            3934 pps
    Bytes              :        222942680        1447768 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
      Low              :                0                0 pps
      Medium-low       :                0                0 pps
      Medium-high      :                0                0 pps
      High              :                0                0 pps
    RED-dropped bytes  :                0                0 bps
      Low              :                0                0 bps
      Medium-low       :                0                0 bps
      Medium-high      :                0                0 bps
      High              :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets           :                0                0 pps
    Bytes              :                0                0 bps
  Transmitted:
    Packets           :                0                0 pps
    Bytes              :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
      Low              :                0                0 pps
      Medium-low       :                0                0 pps
      Medium-high      :                0                0 pps
      High              :                0                0 pps
    RED-dropped bytes  :                0                0 bps
      Low              :                0                0 bps
      Medium-low       :                0                0 bps
      Medium-high      :                0                0 bps
      High              :                0                0 bps

```

show interfaces queue (Gigabit Ethernet IQ2 PIC)

user@host> **show interfaces queue ge-7/1/3**

Physical interface: ge-7/1/3, Enabled, Physical link is Up

Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in use Ingress queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	418390039	10 pps
Bytes	:	38910269752	7440 bps

Transmitted:

Packets	:	418390039	10 pps
Bytes	:	38910269752	7440 bps
Tail-dropped packets : Not Available			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 3, Forwarding classes: network-control

Queued:

Packets	:	7055	1 pps
Bytes	:	451552	512 bps

Transmitted:

Packets	:	7055	1 pps
Bytes	:	451552	512 bps
Tail-dropped packets : Not Available			
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	1031	0 pps
Bytes	:	143292	0 bps

Transmitted:

Packets	:	1031	0 pps
Bytes	:	143292	0 bps
Tail-dropped packets : Not Available			
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 3, Forwarding classes: network-control

Queued:

Packets	:	77009	11 pps
Bytes	:	6894286	7888 bps

Transmitted:

Packets	:	77009	11 pps
Bytes	:	6894286	7888 bps
Tail-dropped packets : Not Available			

RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps

Transmitted:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding


```

Queued:
  Packets      :                0                0 pps
  Bytes       :                0                0 bps
Transmitted:
  Packets      :                0                0 pps
  Bytes       :                0                0 bps
  Tail-dropped packets :                0                0 pps
  RED-dropped packets :                0                0 pps
    Low, non-TCP      :                0                0 pps
    Low, TCP          :                0                0 pps
    High, non-TCP     :                0                0 pps
    High, TCP         :                0                0 pps
  RED-dropped bytes :                0                0 bps
    Low, non-TCP      :                0                0 bps
    Low, TCP          :                0                0 bps
    High, non-TCP     :                0                0 bps
    High, TCP         :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :                94386            12 pps
  Bytes       :            13756799            9568 bps
Transmitted:
  Packets      :                94386            12 pps
  Bytes       :            13756799            9568 bps
  Tail-dropped packets :                0                0 pps
  RED-dropped packets :                0                0 pps
    Low, non-TCP      :                0                0 pps
    Low, TCP          :                0                0 pps
    High, non-TCP     :                0                0 pps
    High, TCP         :                0                0 pps
  RED-dropped bytes :                0                0 bps
    Low, non-TCP      :                0                0 bps
    Low, TCP          :                0                0 bps
    High, non-TCP     :                0                0 bps
    High, TCP         :                0                0 bps

```

show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

user@host> **show interfaces queue ge-6/2/0 both-ingress-egress**

```

Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use

```

```

Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          : Not Available
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                254                0 pps
    Bytes            :            16274                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          : Not Available
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets          : Not Available
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets          : Not Available
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```

```

Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                3                0 pps
  Bytes           :               126                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : Not Available
  Bytes           :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes           :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :                0                0 pps
  RED-dropped bytes  :                0                0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets          :               80564692            0 pps

```

```

    Bytes          :          3383717100          0 bps
Transmitted:
    Packets        :          80564692          0 pps
    Bytes          :          3383717100          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
    Packets        :          80564685          0 pps
    Bytes          :          3383716770          0 bps
Transmitted:
    Packets        :          80564685          0 pps
    Bytes          :          3383716770          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps
Transmitted:
    Packets        :          0          0 pps
    Bytes          :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
    Packets        :          9397          0 pps
    Bytes          :          3809052          232 bps
Transmitted:
    Packets        :          9397          0 pps
    Bytes          :          3809052          232 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps

```

show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

user@host> **show interfaces queue ge-6/2/0 ingress**

Physical interface: ge-6/2/0, Enabled, Physical link is Up

Interface index: 175, SNMP ifIndex: 121

Forwarding classes: 8 supported, 4 in use

Ingress queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	Not Available	
Bytes	:	0	0 bps

Transmitted:

Packets	:	288	0 pps
Bytes	:	18450	0 bps
Tail-dropped packets	:	Not Available	
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	Not Available	
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	Not Available	
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	Not Available	
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	Not Available	
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 3, Forwarding classes: network-control

Queued:

Packets	:	Not Available	
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	Not Available	
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

user@host> **show interfaces queue ge-6/2/0 egress**

```
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          : Not Available
    Bytes            :                      0          0 bps
  Transmitted:
    Packets          :                      3          0 pps
    Bytes            :                   126          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes   :                      0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          : Not Available
    Bytes            :                      0          0 bps
  Transmitted:
    Packets          :                      0          0 pps
    Bytes            :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes   :                      0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets          : Not Available
    Bytes            :                      0          0 bps
  Transmitted:
    Packets          :                      0          0 pps
    Bytes            :                      0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                      0          0 pps
    RED-dropped bytes   :                      0          0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets          : Not Available
    Bytes            :                      0          0 bps
  Transmitted:
    Packets          :                      0          0 pps
```

```

    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                0                0 pps
    RED-dropped bytes    :                0                0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :            80564692            0 pps
    Bytes                :            3383717100          0 bps
  Transmitted:
    Packets              :            80564692            0 pps
    Bytes                :            3383717100          0 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets  :                0                0 pps
    RED-dropped bytes    :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :            80564685            0 pps
    Bytes                :            3383716770          0 bps
  Transmitted:
    Packets              :            80564685            0 pps
    Bytes                :            3383716770          0 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets  :                0                0 pps
    RED-dropped bytes    :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets  :                0                0 pps
    RED-dropped bytes    :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              :                9538            0 pps
    Bytes                :            3819840            0 bps
  Transmitted:
    Packets              :                9538            0 pps
    Bytes                :            3819840            0 bps
    Tail-dropped packets :                0                0 pps

```

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

user@host> show interfaces queue ge-2/2/9 remaining-traffic

```

Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 110208969 472875 pps
    Bytes : 5951284434 204282000 bps
  Transmitted:
    Packets : 110208969 472875 pps
    Bytes : 5951284434 204282000 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps

```



```

    Low                :                0                0 bps
    Medium-low         :                0                0 bps
    Medium-high        :                0                0 bps
    High               :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets           :                0                0 pps
    Bytes             :                0                0 bps
  Transmitted:
    Packets           :                0                0 pps
    Bytes             :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
      Low             :                0                0 pps
      Medium-low      :                0                0 pps
      Medium-high     :                0                0 pps
      High            :                0                0 pps
    RED-dropped bytes :                0                0 bps
      Low             :                0                0 bps
      Medium-low      :                0                0 bps
      Medium-high     :                0                0 bps
      High            :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets           :                0                0 pps
    Bytes             :                0                0 bps
  Transmitted:
    Packets           :                0                0 pps
    Bytes             :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
      Low             :                0                0 pps
      Medium-low      :                0                0 pps
      Medium-high     :                0                0 pps
      High            :                0                0 pps
    RED-dropped bytes :                0                0 bps
      Low             :                0                0 bps
      Medium-low      :                0                0 bps
      Medium-high     :                0                0 bps
      High            :                0                0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:

```

```

Packets      :          109355853          471736 pps
Bytes        :          7436199152        256627968 bps
Transmitted:
Packets      :          109355852          471736 pps
Bytes        :          7436198640        256627968 bps
Tail-dropped packets : Not Available
RED-dropped packets :          0          0 pps
  Low         :          0          0 pps
  Medium-low  :          0          0 pps
  Medium-high :          0          0 pps
  High        :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low         :          0          0 bps
  Medium-low  :          0          0 bps
  Medium-high :          0          0 bps
  High        :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets      :          0          0 pps
Bytes        :          0          0 bps
Transmitted:
Packets      :          0          0 pps
Bytes        :          0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :          0          0 pps
  Low         :          0          0 pps
  Medium-low  :          0          0 pps
  Medium-high :          0          0 pps
  High        :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low         :          0          0 bps
  Medium-low  :          0          0 bps
  Medium-high :          0          0 bps
  High        :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets      :          0          0 pps
Bytes        :          0          0 bps
Transmitted:
Packets      :          0          0 pps
Bytes        :          0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :          0          0 pps
  Low         :          0          0 pps

```

```

Medium-low      :                0                0 pps
Medium-high     :                0                0 pps
High            :                0                0 pps
RED-dropped bytes :                0                0 bps
Low             :                0                0 bps
Medium-low      :                0                0 bps
Medium-high     :                0                0 bps
High            :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets         :                0                0 pps
Bytes           :                0                0 bps
Transmitted:
Packets         :                0                0 pps
Bytes           :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
Low             :                0                0 pps
Medium-low      :                0                0 pps
Medium-high     :                0                0 pps
High            :                0                0 pps
RED-dropped bytes :                0                0 bps
Low             :                0                0 bps
Medium-low      :                0                0 bps
Medium-high     :                0                0 bps
High            :                0                0 bps

```

show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

user@host> show interfaces queue t3-1/1/0:7

```

Physical interface: t3-1/1/0:7, Enabled, Physical link is Up
Interface index: 192, SNMP ifIndex: 1948
Description: full T3 interface connect to 6cel3 t3-3/1/0:7 for FR testing - Lam

Forwarding classes: 16 supported, 9 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: DEFAULT
Queued:
Packets         :                214886                13449 pps
Bytes           :                9884756                5164536 bps
Transmitted:
Packets         :                214886                13449 pps

```

Bytes	:	9884756	5164536 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

Transmitted:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High	:	0	0 bps
Queue: 2, Forwarding classes: PRIVATE			
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps
Queue: 3, Forwarding classes: CONTROL			
Queued:			
Packets	:	32843	0 pps
Bytes	:	2641754	56 bps
Transmitted:			
Packets	:	32843	0 pps
Bytes	:	2641754	56 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps
Queue: 4, Forwarding classes: CLASS_B_OUTPUT			
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps

Tail-dropped packets :	0	0 pps
RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps
Medium-low :	0	0 bps
Medium-high :	0	0 bps
High :	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets :	0	0 pps
Bytes :	0	0 bps

Transmitted:

Packets :	0	0 pps
Bytes :	0	0 bps
Tail-dropped packets :	0	0 pps
RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps
Medium-low :	0	0 bps
Medium-high :	0	0 bps
High :	0	0 bps

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets :	0	0 pps
Bytes :	0	0 bps

Transmitted:

Packets :	0	0 pps
Bytes :	0	0 bps
Tail-dropped packets :	0	0 pps
RED-dropped packets :	0	0 pps
Low :	0	0 pps
Medium-low :	0	0 pps
Medium-high :	0	0 pps
High :	0	0 pps
RED-dropped bytes :	0	0 bps
Low :	0	0 bps

```

Medium-low      :                0                0 bps
Medium-high     :                0                0 bps
High            :                0                0 bps
Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS
Queued:
Packets         :                0                0 pps
Bytes           :                0                0 bps
Transmitted:
Packets         :                0                0 pps
Bytes           :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets :                0                0 pps
Low             :                0                0 pps
Medium-low      :                0                0 pps
Medium-high     :                0                0 pps
High            :                0                0 pps
RED-dropped bytes :                0                0 bps
Low             :                0                0 bps
Medium-low      :                0                0 bps
Medium-high     :                0                0 bps
High            :                0                0 bps

```

show interfaces queue (QFX Series)

user@switch> **show interfaces queue xe-0/0/15**

```

Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets         :                0                0 pps
Bytes           :                0                0 bps
Transmitted:
Packets         :                0                0 pps
Bytes           :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped packets:                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 3, Forwarding classes: fcoe
Queued:
Packets         :                0                0 pps

```

```

        Bytes                :                0                0 bps
    Transmitted:
        Packets                :                0                0 pps
        Bytes                  :                0                0 bps
        Tail-dropped packets : Not Available
        Total-dropped packets:                0                0 pps
        Total-dropped bytes  :                0                0 bps
Queue: 4, Forwarding classes: no-loss
    Queued:
        Packets                :                0                0 pps
        Bytes                  :                0                0 bps
    Transmitted:
        Packets                :                0                0 pps
        Bytes                  :                0                0 bps
        Tail-dropped packets : Not Available
        Total-dropped packets:                0                0 pps
        Total-dropped bytes  :                0                0 bps
Queue: 7, Forwarding classes: network-control
    Queued:
        Packets                :                0                0 pps
        Bytes                  :                0                0 bps
    Transmitted:
        Packets                :                0                0 pps
        Bytes                  :                0                0 bps
        Tail-dropped packets : Not Available
        Total-dropped packets:                0                0 pps
        Total-dropped bytes  :                0                0 bps
Queue: 8, Forwarding classes: mcast
    Queued:
        Packets                :                0                0 pps
        Bytes                  :                0                0 bps
    Transmitted:
        Packets                :                0                0 pps
        Bytes                  :                0                0 bps
        Tail-dropped packets : Not Available
        Total-dropped packets:                0                0 pps
        Total-dropped bytes  :                0                0 bps

```

show interfaces queue l2-statistics (lsq interface)

user@switch> **show interfaces queue lsq-2/2/0.2 l2-statistics**

Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)

Forwarding classes: 16 supported, 4 in use

Egress queues: 8 supported, 4 in use

Burst size: 0

Queue: 0, Forwarding classes: be

Queued:

Packets	:	1	0 pps
Bytes	:	1001	0 bps

Transmitted:

Packets	:	5	0 pps
Bytes	:	1062	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 1, Forwarding classes: ef

Queued:

Packets	:	1	0 pps
Bytes	:	1500	0 bps

Transmitted:

Packets	:	6	0 pps
Bytes	:	1573	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: af

Queued:

Packets	:	1	0 pps
Bytes	:	512	0 bps

Transmitted:

Packets	:	3	0 pps
Bytes	:	549	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
RED-dropped bytes	:	0	0 bps

Queue: 3, Forwarding classes: nc

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps

```

RED-dropped bytes      :                0                0 bps
=====

```

show interfaces queue lsq (lsq-afd)

```
user@switch> show interfaces queue lsq-1/0/0
```

```

Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :                55576                1206 pps
    Bytes        :            29622008            5145472 bps
  Transmitted:
    Packets      :                55576                1206 pps
    Bytes        :            29622008            5145472 bps
    Tail-dropped packets :                0                0 pps
    RL-dropped packets  :                0                0 pps
    RL-dropped bytes    :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low              :                0                0 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                0                0 pps
    RED-dropped bytes   :                0                0 bps
    Low                :                0                0 bps
    Medium-low         :                0                0 bps
    Medium-high        :                0                0 bps
    High               :                0                0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RL-dropped packets  :                0                0 pps
    RL-dropped bytes    :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low              :                0                0 pps

```

```

    Medium-low      :                0                0 pps
    Medium-high     :                0                0 pps
    High            :                0                0 pps
    RED-dropped bytes :                0                0 bps
    Low             :                0                0 bps
    Medium-low      :                0                0 bps
    Medium-high     :                0                0 bps
    High            :                0                0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
  Transmitted:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RL-dropped packets :                0                0 pps
    RL-dropped bytes  :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low             :                0                0 pps
    Medium-low      :                0                0 pps
    Medium-high     :                0                0 pps
    High            :                0                0 pps
    RED-dropped bytes :                0                0 bps
    Low             :                0                0 bps
    Medium-low      :                0                0 bps
    Medium-high     :                0                0 bps
    High            :                0                0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets        :            22231            482 pps
    Bytes          :       11849123       2057600 bps
  Transmitted:
    Packets        :            22231            482 pps
    Bytes          :       11849123       2057600 bps
    Tail-dropped packets :                0                0 pps
    RL-dropped packets :                0                0 pps
    RL-dropped bytes  :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low             :                0                0 pps
    Medium-low      :                0                0 pps
    Medium-high     :                0                0 pps
    High            :                0                0 pps
    RED-dropped bytes :                0                0 bps

```

Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

show interfaces queue (Aggregated Ethernet on a MX series Router)

user@host> show interfaces queue ae0 remaining-traffic

```
Physical interface: ae0      , Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 543
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :                16          0 pps
    Bytes        :               1896          0 bps
  Transmitted:
    Packets      :                16          0 pps
    Bytes        :               1896          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-depth bytes :
    Average        :                0
    Current         :                0
    Peak           :                0
    Maximum         :             119013376
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
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-depth bytes :
  Average        :                0
  Current        :                0
  Peak           :                0
  Maximum        :                32768
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets        :                0                0 pps
  Bytes          :                0                0 bps
Transmitted:
  Packets        :                0                0 pps
  Bytes          :                0                0 bps
Tail-dropped packets :                0                0 pps
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-depth bytes :
  Average        :                0
  Current        :                0
  Peak           :                0

```



```

Maximum          :          32768
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
  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-depth bytes  :
  Average          :          0
  Current          :          0
  Peak             :          0
  Maximum          :        6258688

```

show interfaces queue ge-0/0/0 (EX2200 Switch)

```
user@switch> show interfaces queue ge-0/0/0
```

```

Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 130, SNMP ifIndex: 501
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets          :          0
    Bytes            :          0
    Tail-dropped packets :          0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:

```

```

Packets          :                0
Bytes            :                0
Tail-dropped packets :                0
Queue: 5, Forwarding classes: expedited-forwarding
Queued:
Transmitted:
Packets          :                0
Bytes            :                0
Tail-dropped packets :                0
Queue: 7, Forwarding classes: network-control
Queued:
Transmitted:
Packets          :                0
Bytes            :                0
Tail-dropped packets :                0

```

show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch)

user@switch> show interfaces queue xe-6/0/39

```

Physical interface: xe-6/0/39, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 1641
Forwarding classes: 16 supported, 7 in use
Ingress queues: 1 supported, 1 in use
  Transmitted:
    Packets          :          337069086018
    Bytes            :          43144843010304
    Tail-dropped packets :          8003867575
PFE chassis queues: 1 supported, 1 in use
  Transmitted:
    Packets          :                0
    Bytes            :                0
    Tail-dropped packets :                0
Forwarding classes: 16 supported, 7 in use
Egress queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Transmitted:
  Packets          :          334481399932
  Bytes            :          44151544791024
  Tail-dropped packets :                0
Queue: 1, Forwarding classes: assured-forwarding
Queued:

```

```

Transmitted:
  Packets          :          0
  Bytes           :          0
  Tail-dropped packets :          0
Queue: 2, Forwarding classes: mcast-be
Queued:
Transmitted:
  Packets          :      274948977
  Bytes           :      36293264964
  Tail-dropped packets :          0
Queue: 4, Forwarding classes: mcast-ef
Queued:
Transmitted:
  Packets          :          0
  Bytes           :          0
  Tail-dropped packets :          0
Queue: 5, Forwarding classes: expedited-forwarding
Queued:
Transmitted:
  Packets          :          0
  Bytes           :          0
  Tail-dropped packets :          0
Queue: 6, Forwarding classes: mcast-af
Queued:
Transmitted:
  Packets          :          0
  Bytes           :          0
  Tail-dropped packets :          0
Queue: 7, Forwarding classes: network-control
Queued:
Transmitted:
  Packets          :      46714
  Bytes           :      6901326
  Tail-dropped packets :          0

Packet Forwarding Engine Chassis Queues:
Queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Transmitted:
  Packets          :      739338141426
  Bytes           :      94635282101928
  Tail-dropped packets :          0
  RED-dropped packets :      5606426444

```

```

    Low                :                5606426444
    High               :                0
    RED-dropped bytes  :                683262846464
    Low                :                683262846464
    High               :                0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets           :                0
    Bytes             :                0
    Tail-dropped packets :                0
    RED-dropped packets :                0
    Low               :                0
    High              :                0
    RED-dropped bytes  :                0
    Low               :                0
    High              :                0
Queue: 2, Forwarding classes: mcast-be
  Queued:
  Transmitted:
    Packets           :                0
    Bytes             :                0
    Tail-dropped packets :                0
    RED-dropped packets :                0
    Low               :                0
    High              :                0
    RED-dropped bytes  :                0
    Low               :                0
    High              :                0
Queue: 4, Forwarding classes: mcast-ef
  Queued:
  Transmitted:
    Packets           :                0
    Bytes             :                0
    Tail-dropped packets :                0
    RED-dropped packets :                0
    Low               :                0
    High              :                0
    RED-dropped bytes  :                0
    Low               :                0
    High              :                0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:

```

```

Packets          :          0
Bytes            :          0
Tail-dropped packets :          0
RED-dropped packets :          0
  Low            :          0
  High           :          0
RED-dropped bytes :          0
  Low            :          0
  High           :          0
Queue: 6, Forwarding classes: mcast-af
  Queued:
    Transmitted:
      Packets          :          0
      Bytes            :          0
      Tail-dropped packets :          0
      RED-dropped packets :          0
        Low            :          0
        High           :          0
      RED-dropped bytes :          0
        Low            :          0
        High           :          0
Queue: 7, Forwarding classes: network-control
  Queued:
    Transmitted:
      Packets          :          97990
      Bytes            :         14987506
      Tail-dropped packets :          0
      RED-dropped packets :          0
        Low            :          0
        High           :          0
      RED-dropped bytes :          0
        Low            :          0
        High           :          0

```

show interfaces queue xe-0/0/2 buffer-occupancy (QFX5000 Series switch)

user@switch> show interfaces queue xe-0/0/2 buffer-occupancy

```

Physical interface: xe-0/0/2, Enabled, Physical link is Up
  Interface index: 689, SNMP ifIndex: 514
Forwarding classes: 12 supported, 5 in use
Egress queues: 12 supported, 5 in use
  Queue: 0, Forwarding classes: fc0

```

```
Queue-depth bytes  :  
Peak               : 1786720  
Queue: 3, Forwarding classes: fcoe  
Queue-depth bytes  :  
Peak               : 0  
Queue: 4, Forwarding classes: no-loss  
Queue-depth bytes  :  
Peak               : 0  
Queue: 7, Forwarding classes: network-control  
Queue-depth bytes  :  
Peak               : 416  
Queue: 8, Forwarding classes: mcast  
Queue-depth bytes  :  
Peak               : 0
```

show interfaces queue fabric

Syntax

```
show interfaces queue fabric
<egress>
<forwarding-class forwarding-class>
<interface-name interface-name>
```

Release Information

Command introduced in Junos OS Release 12.3 for the QFX Series.

Description

Display class-of-service (CoS) queue information for the fabric interfaces that are configured between Node devices and Interconnect devices.

Options

none—Show detailed CoS queue statistics for all physical interfaces.

egress—(Optional) Display egress queue statistics.

forwarding-class *forwarding-class*—(Optional) Forwarding class name for this queue. Show detailed CoS statistics for the queue associated with the specified forwarding class.

interface-name *interface-name*—(Optional) Show detailed CoS queue statistics for the specified interface.

Required Privilege Level

view

RELATED DOCUMENTATION

[show interfaces fabric | 1047](#)

List of Sample Output

[show interfaces queue fabric on page 1185](#)

[show interfaces queue fabric egress on page 1197](#)

[show interfaces queue fabric interface-name egress on page 1210](#)

[show interfaces queue fabric interface-name egress forwarding-class forwarding-class-name on page 1211](#)

Output Fields

[Table 90 on page 1132](#) lists the output fields for the **show interfaces queue fabric** command. Output fields are listed in the approximate order in which they appear.

Table 92: show interfaces queue fabric Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	<p>State of the interface. Possible values are:</p> <ul style="list-style-type: none"> • Administratively down, Physical link is Down–The interface is turned off, and the physical link is inoperable. • Administratively down, Physical link is Up–The interface is turned off, but the physical link is operational and can pass packets when it is enabled. • Enabled, Physical link is Down–The interface is turned on, but the physical link is inoperable and cannot pass packets. • Enabled, Physical link is Up–The interface is turned on, and the physical link is operational and can pass packets.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Forwarding classes	Number of forwarding classes supported and in use for the interface.
Egress queues	Number of output queues supported and in use on the specified interface.
Queue	CoS queue number.
Transmitted	<p>Number of packets and bytes transmitted by this queue. Information on transmitted packets and bytes can include:</p> <ul style="list-style-type: none"> • Packets–Number of packets transmitted. • Bytes–Number of bytes transmitted. • Tail-dropped packets–Number of arriving packets dropped because output queue buffers were full. • Total-dropped pkts–Number of transmitted packets dropped. • Total dropped bytes–Number of transmitted bytes dropped.
Queued	<p>Number of packets and bytes queued to this queue.</p> <ul style="list-style-type: none"> • Packets–Number of packets queued. • Bytes–Number of bytes queued.

Sample Output

show interfaces queue fabric

user@switch> show interfaces queue fabric

```
Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
  Interface index: 49178, SNMP ifIndex: 1208484475
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :          62665971                0 pps
    Bytes            :      7770580404                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts  :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts  :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts  :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps

```

```

    Total-dropped bytes      :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
Interface index: 49211, SNMP ifIndex: 1208484377

```

Forwarding classes: 16 supported, 5 in use

Egress queues: 12 supported, 5 in use

Queue: 0, Forwarding classes: fabric_fcset_be

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

Queue: 1, Forwarding classes: fabric_fcset_noloss1

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

Queue: 2, Forwarding classes: fabric_fcset_noloss2

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

Queue: 3, Forwarding classes: fabric_fcset_noloss3

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

Queue: 4, Forwarding classes: fabric_fcset_noloss4

Queued:

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps

```

```

    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
  Interface index: 49212, SNMP ifIndex: 1208484365
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:

```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps

```

```

    Total-dropped bytes :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                0                0 pps

```



```

    Bytes          :                0                0 bps
Transmitted:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts  :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
Transmitted:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts  :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
Transmitted:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts  :                0                0 pps
    Total-dropped bytes :                0                0 bps

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
Interface index: 49213, SNMP ifIndex: 1208484625
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
Queued:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
Transmitted:
    Packets        :                0                0 pps
    Bytes          :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts  :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
Queued:

```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps

```

```

    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps

```

Queue: 10, Forwarding classes: fabric_fcset_mcast3

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

Queue: 11, Forwarding classes: fabric_fcset_mcast4

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up

Interface index: 49161, SNMP ifIndex: 1209008630

Forwarding classes: 16 supported, 5 in use

Egress queues: 12 supported, 5 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

Queue: 3, Forwarding classes: fcoe

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps

```

    Total-dropped bytes :                0                0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps

```

show interfaces queue fabric egress

user@switch> show interfaces queue fabric egress

```

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Up
  Interface index: 49178, SNMP ifIndex: 1208484475
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps

```

```

Transmitted:
  Packets          :          62665971          0 pps
  Bytes            :          7770580404        0 bps
  Tail-dropped packets : Not Available
  Total-dropped pkts   :              0          0 pps
  Total-dropped bytes  :              0          0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
  Transmitted:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :              0          0 pps
    Total-dropped bytes  :              0          0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
  Transmitted:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :              0          0 pps
    Total-dropped bytes  :              0          0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
  Transmitted:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :              0          0 pps
    Total-dropped bytes  :              0          0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
  Transmitted:
    Packets          :              0          0 pps
    Bytes            :              0          0 bps
    Tail-dropped packets : Not Available

```

```

    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:

```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts  :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts  :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts  :                0                0 pps
Total-dropped bytes :                0                0 bps

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Up
Interface index: 49211, SNMP ifIndex: 1208484377
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts  :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1

```



```

Queued:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
  Tail-dropped packets : Not Available
  Total-dropped pkts  :                0                0 pps
  Total-dropped bytes :                0                0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
Queued:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
  Tail-dropped packets : Not Available
  Total-dropped pkts  :                0                0 pps
  Total-dropped bytes :                0                0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
Queued:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
  Tail-dropped packets : Not Available
  Total-dropped pkts  :                0                0 pps
  Total-dropped bytes :                0                0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
Queued:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
Transmitted:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
  Tail-dropped packets : Not Available
  Total-dropped pkts  :                0                0 pps
  Total-dropped bytes :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
Queued:
  Packets          :                0                0 pps
  Bytes            :                0                0 bps
Transmitted:

```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps

```

```

    Total-dropped bytes      :                0                0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Up
  Interface index: 49212, SNMP ifIndex: 1208484365
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available

```

```

    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
  Transmitted:
    Packets                 :                0                0 pps
    Bytes                   :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:

```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps

```

```

    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Up
  Interface index: 49213, SNMP ifIndex: 1208484625
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: fabric_fcset_be
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: fabric_fcset_noloss1
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: fabric_fcset_noloss2
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:

```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 3, Forwarding classes: fabric_fcset_noloss3
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 4, Forwarding classes: fabric_fcset_noloss4
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 5, Forwarding classes: fabric_fcset_noloss5
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps
    Total-dropped bytes :                0                0 bps
Queue: 6, Forwarding classes: fabric_fcset_noloss6
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts :                0                0 pps

```

```

    Total-dropped bytes      :                0                0 bps
Queue: 7, Forwarding classes: fabric_fcset_strict_high
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 8, Forwarding classes: fabric_fcset_mcast1
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 9, Forwarding classes: fabric_fcset_mcast2
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 10, Forwarding classes: fabric_fcset_mcast3
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 11, Forwarding classes: fabric_fcset_mcast4
  Queued:
    Packets                  :                0                0 pps

```



```

    Bytes                :                0                0 bps
Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps

Physical interface: P2659-C:fte-0/1/2, Enabled, Physical link is Up
Interface index: 49161, SNMP ifIndex: 1209008630
Forwarding classes: 16 supported, 5 in use
Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
  Transmitted:
    Packets              :                0                0 pps
    Bytes                :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts   :                0                0 pps
    Total-dropped bytes  :                0                0 bps
Queue: 7, Forwarding classes: network-control
  Queued:

```

```

Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 8, Forwarding classes: mcast
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps

```

show interfaces queue fabric interface-name egress

user@switch> **show interfaces queue fabric BBAK0394:fte-0/1/0 egress**

```

Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
Interface index: 129, SNMP ifIndex: 1091568120 Forwarding classes: 16 supported,
5 in use Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :          74777763341          844587 pps
Bytes            :          9272442654284        837830728 bps
Tail-dropped packets : Not Available
Total-dropped pkts :                0                0 pps
Total-dropped bytes :                0                0 bps
Queue: 3, Forwarding classes: fcoe
Queued:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available

```

```

    Total-dropped pkts      :                0                0 pps
    Total-dropped bytes     :                0                0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
  Transmitted:
    Packets                  :                0                0 pps
    Bytes                    :                0                0 bps
    Tail-dropped packets : Not Available
    Total-dropped pkts     :                0                0 pps
    Total-dropped bytes    :                0                0 bps

```

show interfaces queue fabric interface-name egress forwarding-class forwarding-class-name

user@switch> **show interfaces queue fabric BBAK0394:fte-0/1/0 egress forwarding-class best-effort**

```

Physical interface: BBAK0394:fte-0/1/0, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 1091568120 Forwarding classes: 16 supported,
  5 in use Egress queues: 12 supported, 5 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                  :                0                0 pps

```

Bytes	:	0	0 bps
Transmitted:			
Packets	:	74793424543	844612 pps
Bytes	:	9274384643332	837855936 bps
Tail-dropped packets : Not Available			
Total-dropped pkts	:	0	0 pps
Total-dropped bytes	:	0	0 bps

show interfaces xe

List of Syntax

[Syntax \(QFX Series\) on page 1213](#)

[Syntax \(EX Series\) on page 1213](#)

Syntax (QFX Series)

```
show interfaces device-name:type-fpc/pic/port  
<brief | detail | extensive | terse>  
<descriptions>  
<media>  
<routing-instance (all | instance-name)>  
<snmp-index snmp-index>  
<statistics>
```

Syntax (EX Series)

```
show interfaces xe-fpc/pic/port  
<brief | detail | extensive | terse>  
<media>  
<statistics>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.

NOTE: You must have a transceiver plugged into an SFP+ or an XFP port before information about the interface can be displayed.

NOTE: On an EX Series switch, the traffic statistics for a LAG might vary slightly from the cumulative traffic statistics of the member interfaces of the LAG. This difference is more likely to be seen when the traffic is bursty in nature, and because the statistics are not fetched from the LAG and the members in the same instant. For accurate traffic statistics for a LAG, use the aggregated Ethernet counters.

Options

device-name:*type-fpc/pic/port*—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.

brief | detail | extensive | terse—(Optional) (QFX Series) Display the specified level of output.

descriptions—(Optional) (QFX Series) Display interface description strings.

media—(Optional) (QFX Series) Display media-specific information about network interfaces.

routing-instance (**all** | **instance-name**)—(Optional) (QFX Series) Display the name of an individual routing instance or display all routing instances.

snmp-index *snmp-index*—(Optional) (QFX Series) Display information for the specified SNMP index of the interface.

statistics—(Optional) (QFX Series) Display static interface statistics.

xe-fpc/pic/port —(EX Series) Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) (EX Series) Display the specified level of output.

media—(Optional) (EX Series) Display media-specific information about network interfaces. For 10-Gigabit Ethernet interfaces, using the media option does not provide you with new or additional information. The output is the same as when the media option is not used.

statistics—(Optional) (EX Series) Display static interface statistics. For 10-Gigabit Ethernet interfaces, using the statistics option does not provide you with new or additional information. The output is the same as when the statistics option is not used.

Required Privilege Level

view

RELATED DOCUMENTATION

[Monitoring Interface Status and Traffic](#) | 392

Monitoring Interface Status and Traffic

[Troubleshooting Network Interfaces | 399](#)

Troubleshooting an Aggregated Ethernet Interface

Junos OS Network Interfaces Library for Routing Devices

Troubleshooting Network Interfaces on EX3200 Switches

Troubleshooting Network Interfaces on EX4200 Switches

[Troubleshooting an Aggregated Ethernet Interface | 279](#)

[Junos OS Ethernet Interfaces Configuration Guide](#)

List of Sample Output

[show interfaces on page 1225](#)

[show interfaces \(Asymmetric Flow Control\) on page 1225](#)

[show interfaces brief on page 1226](#)

[show interfaces detail on page 1226](#)

[show interfaces detail \(Asymmetric Flow Control\) on page 1228](#)

[show interfaces extensive on page 1230](#)

[show interfaces extensive \(Asymmetric Flow Control\) on page 1232](#)

[show interfaces terse on page 1235](#)

[show interfaces \(QFabric System\) on page 1235](#)

Output Fields

Table 83 on page 1048 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

Table 93: show interfaces xe Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels

Table 93: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex .	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled . NOTE: This field is only displayed if asymmetric flow control is not configured.	All levels
Configured-flow-control	Configured flow control for the interface transmit buffers (tx-buffers) and receive buffers (rx-buffers): <ul style="list-style-type: none"> • tx-buffers—On if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer. Off if the interface is not configured to respond to received PAUSE messages. • rx-buffers—On if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer. Off if the interface is not configured to generate and send PAUSE messages. NOTE: This field is only displayed if asymmetric flow control is configured.	All levels
Auto-negotiation	Autonegotiation status: Enabled or Disabled .	All levels

Table 93: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Remote-fault	Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
Wavelength	Configured wavelength, in nanometers (nm).	All levels
Frequency	Frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour: :minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 93: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 93: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 93: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Queue Number	The CoS queue number and the forwarding classes mapped to the queue number. The Mapped forwarding class column lists the forwarding classes mapped to each CoS queue.	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link . <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
PCS statistics	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	detail extensive

Table 93: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of packets that exceeds the configured MTU. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
Filter statistics	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.	extensive

Table 93: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are enable or disable. • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive

Table 93: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	detail extensive none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	detail extensive
IPv6 transit statistics	If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	extensive

Table 93: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to and from the switch.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Input Filters	Names of any input filters applied to this interface.	detail extensive
Output Filters	Names of any output filters applied to this interface.	detail extensive
Flags	Information about protocol family flags. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	detail extensive
Addresses, Flags	Information about the address flags.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces

user@switch> **show interfaces xe-0/0/1**

```
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues       : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped    : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode
```

show interfaces (Asymmetric Flow Control)

user@switch> **show interfaces xe-0/0/1**

```
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
```

```

CoS queues      : 12 supported, 12 maximum usable queues
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped    : 2011-06-01 00:42:03 PDT (00:02:42 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
Active alarms   : None
Active defects  : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode

```

show interfaces brief

user@switch> **show interfaces xe-0/0/1 brief**

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags        : None

Logical interface xe-0/0/1.0
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch

```

show interfaces detail

user@switch> **show interfaces xe-0/0/1 detail**

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

```

```

Link flags      : None
CoS queues     : 12 supported, 12 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped   : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
  Input  bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input  packets:                0                0 pps
  Output packets:                0                0 pps
IPv6 transit statistics:
  Input  bytes :                0
  Output bytes :                0
  Input  packets:                0
  Output packets:                0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort      0                0                0
  1 fc7              0                0                0
  2 no-loss          0                0                0
  3 fcoe             0                0                0
  4 fc4              0                0                0
  5 fc5              0                0                0
  6 fc6              0                0                0
  7 network-cont     0                0                0
  8 mcast            0                0                0
Queue number:      Mapped forwarding classes
  0                best-effort
  1                fc7
  2                no-loss
  3                fcoe
  4                fc4
  5                fc5
  6                fc6
  7                network-control
  8                mcast
Active alarms  : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes :                0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces detail (Asymmetric Flow Control)

user@switch> show interfaces xe-0/0/1 detail

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Egress queues: 12 supported, 9 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 fc7	0	0	0
2 no-loss	0	0	0
3 fcoe	0	0	0
4 fc4	0	0	0
5 fc5	0	0	0
6 fc6	0	0	0
7 network-cont	0	0	0
8 mcast	0	0	0

Queue number:	Mapped forwarding classes
0	best-effort
1	fc7
2	no-loss
3	fcoe
4	fc4
5	fc5
6	fc6
7	network-control
8	mcast

Active alarms : None

Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Local statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Transit statistics:

Input bytes :	0	0 bps
Output bytes :	0	0 bps
Input packets:	0	0 pps
Output packets:	0	0 pps

Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0

Flags: Trunk-Mode

show interfaces extensive

```
user@switch> show interfaces xe-0/0/1 extensive
```

```
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues       : 12 supported, 12 maximum usable queues
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped    : 2011-06-01 00:42:03 PDT (00:03:08 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
    Input bytes   : 0 0 bps
    Output bytes  : 0 0 bps
    Input packets : 0 0 pps
    Output packets: 0 0 pps
  IPv6 transit statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   0 0 0
    1 fc7           0 0 0
    2 no-loss       0 0 0
    3 fcoe          0 0 0
    4 fc4           0 0 0
    5 fc5           0 0 0
    6 fc6           0 0 0
    7 network-cont  0 0 0
```

```

      8 mcast                                0                                0                                0
Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast

Active alarms   : None
Active defects  : None

MAC statistics:                                Receive                                Transmit
Total octets                                0                                0
Total packets                                0                                0
Unicast packets                             0                                0
Broadcast packets                           0                                0
Multicast packets                           0                                0
CRC/Align errors                             0                                0
FIFO errors                                 0                                0
MAC control frames                           0                                0
MAC pause frames                             0                                0
Oversized frames                             0
Jabber frames                               0
Fragment frames                             0
VLAN tagged frames                           0
Code violations                             0

MAC Priority Flow Control Statistics:
Priority : 0                                0                                0
Priority : 1                                0                                0
Priority : 2                                0                                0
Priority : 3                                0                                0
Priority : 4                                0                                0
Priority : 5                                0                                0
Priority : 6                                0                                0
Priority : 7                                0                                0

Filter statistics:
Input packet count                           0
Input packet rejects                         0
Input DA rejects                             0
Input SA rejects                             0
Output packet count                           0
Output packet pad count                       0

```

```

Output packet error count                                0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
    0 best-effort          75      7500000000    75      0      low
none
    7 network-control      5       500000000     5      0      low
none
    8 mcast                20      2000000000    20      0      low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input  bytes :          0
  Output bytes :          0
  Input  packets:          0
  Output packets:          0
Local statistics:
  Input  bytes :          0
  Output bytes :          0
  Input  packets:          0
  Output packets:          0
Transit statistics:
  Input  bytes :          0      0 bps
  Output bytes :          0      0 bps
  Input  packets:          0      0 pps
  Output packets:          0      0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces extensive (Asymmetric Flow Control)

user@switch> show interfaces xe-0/0/1 extensive

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:

```



```

Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues       : 12 supported, 12 maximum usable queues
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped    : 2011-06-01 00:42:03 PDT (00:03:08 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
  IPv6 transit statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort    0                0                0
    1 fc7            0                0                0
    2 no-loss        0                0                0
    3 fcoe           0                0                0
    4 fc4            0                0                0
    5 fc5            0                0                0
    6 fc6            0                0                0
    7 network-cont   0                0                0
    8 mcast          0                0                0
  Queue number:
    Mapped forwarding classes
    0 best-effort
    1 fc7
    2 no-loss
    3 fcoe
    4 fc4

```

```

5          fc5
6          fc6
7          network-control
8          mcast
Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive          Transmit
Total packets         0          0
Unicast packets       0          0
Broadcast packets     0          0
Multicast packets     0          0
CRC/Align errors      0          0
FIFO errors           0          0
MAC control frames    0          0
MAC pause frames      0          0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations       0
MAC Priority Flow Control Statistics:
Priority : 0          0          0
Priority : 1          0          0
Priority : 2          0          0
Priority : 3          0          0
Priority : 4          0          0
Priority : 5          0          0
Priority : 6          0          0
Priority : 7          0          0
Filter statistics:
Input packet count    0
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue    Bandwidth    Buffer Priority    Limit

```

	%	bps	%	usec		
0 best-effort	75	7500000000	75	0	low	none
7 network-control	5	500000000	5	0	low	none
8 mcast	20	2000000000	20	0	low	none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Transit statistics:

Input bytes : 0 0 bps

Output bytes : 0 0 bps

Input packets: 0 0 pps

Output packets: 0 0 pps

Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0

Flags: Trunk-Mode

show interfaces terse

user@switch> show interfaces xe-0/0/1 terse

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/1	up	up			
xe-0/0/1.0	up	up	eth-switch		

show interfaces (QFabric System)

user@switch> show interfaces node1:xe-0/0/0

Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down

Interface index: 129, SNMP ifIndex: 2884086

Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU

Error: None, MAC-REWRITE Error: None,

Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```
Interface flags: Internal: 0x4000
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped    : Never
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
```

show interfaces xle

Syntax

```
show interfaces device-name:type-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced in Junos OS Release 11.1 for the QFX Series.

Description

Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.

Options

device-name:type-fpc/pic/port—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.

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.

routing-instance (all | *instance-name*)—(Optional) Display the name of an individual routing instance or display all routing instances.

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

[Monitoring Interface Status and Traffic | 392](#)

[Troubleshooting Network Interfaces | 399](#)

List of Sample Output[show interfaces on page 1248](#)[show interfaces \(Asymmetric Flow Control\) on page 1248](#)[show interfaces brief on page 1249](#)[show interfaces detail on page 1249](#)[show interfaces detail \(Asymmetric Flow Control\) on page 1251](#)[show interfaces extensive on page 1253](#)[show interfaces extensive \(Asymmetric Flow Control\) on page 1255](#)[show interfaces terse on page 1258](#)[show interfaces \(QFabric System\) on page 1258](#)**Output Fields**

[Table 83 on page 1048](#) lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

Table 94: show interfaces xe Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex .	All levels

Table 94: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled . NOTE: This field is only displayed if asymmetric flow control is not configured.	All levels
Configured-flow-control	Configured flow control for the interface transmit buffers (tx-buffers) and receive buffers (rx-buffers): <ul style="list-style-type: none"> • tx-buffers—On if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer. Off if the interface is not configured to respond to received PAUSE messages. • rx-buffers—On if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer. Off if the interface is not configured to generate and send PAUSE messages. NOTE: This field is only displayed if asymmetric flow control is configured.	All levels
Auto-negotiation	Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels

Table 94: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Wavelength	Configured wavelength, in nanometers (nm).	All levels
Frequency	Frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour: :minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail extensive

Table 94: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 94: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive

Table 94: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Queue Number	The CoS queue number and the forwarding classes mapped to the queue number. The Mapped forwarding class column lists the forwarding classes mapped to each CoS queue.	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link . <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
PCS statistics	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	detail extensive

Table 94: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of packets that exceeds the configured MTU. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
Filter statistics	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.	extensive

Table 94: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are enable or disable. • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive

Table 94: show interfaces xe Output Fields (continued)

Field Name	Field Description	Level of Output
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	detail extensive none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	detail extensive
IPv6 transit statistics	If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	extensive

Table 94: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to and from the switch.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Input Filters	Names of any input filters applied to this interface.	detail extensive
Output Filters	Names of any output filters applied to this interface.	detail extensive
Flags	Information about protocol family flags. If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	detail extensive
Addresses, Flags	Information about the address flags.	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag.	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces

user@switch> show interfaces xe-0/0/1

```
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues       : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped    : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)
  Active alarms    : None
  Active defects   : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode
```

show interfaces (Asymmetric Flow Control)

user@switch> show interfaces xe-0/0/1

```
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
```



```

CoS queues      : 12 supported, 12 maximum usable queues
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped    : 2011-06-01 00:42:03 PDT (00:02:42 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
Active alarms   : None
Active defects  : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode

```

show interfaces brief

user@switch> **show interfaces xe-0/0/1 brief**

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags        : None

Logical interface xe-0/0/1.0
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch

```

show interfaces detail

user@switch> **show interfaces xe-0/0/1 detail**

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

```

```

Link flags      : None
CoS queues      : 12 supported, 12 maximum usable queues
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped    : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
  Input  bytes :                0                0 bps
  Output bytes :                0                0 bps
  Input  packets:                0                0 pps
  Output packets:                0                0 pps
IPv6 transit statistics:
  Input  bytes :                0
  Output bytes :                0
  Input  packets:                0
  Output packets:                0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets
  0 best-effort      0                0                0
  1 fc7              0                0                0
  2 no-loss          0                0                0
  3 fcoe             0                0                0
  4 fc4              0                0                0
  5 fc5              0                0                0
  6 fc6              0                0                0
  7 network-cont     0                0                0
  8 mcast            0                0                0
Queue number:      Mapped forwarding classes
  0                best-effort
  1                fc7
  2                no-loss
  3                fcoe
  4                fc4
  5                fc5
  6                fc6
  7                network-control
  8                mcast
Active alarms      : None
Active defects     : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes :                0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces detail (Asymmetric Flow Control)

user@switch> **show interfaces xe-0/0/1 detail**

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

```

Egress queues: 12 supported, 9 in use

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 fc7	0	0	0
2 no-loss	0	0	0
3 fcoe	0	0	0
4 fc4	0	0	0
5 fc5	0	0	0
6 fc6	0	0	0
7 network-cont	0	0	0
8 mcast	0	0	0

Queue number:	Mapped forwarding classes
0	best-effort
1	fc7
2	no-loss
3	fcoe
4	fc4
5	fc5
6	fc6
7	network-control
8	mcast

Active alarms : None

Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Local statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Transit statistics:

Input bytes :	0	0 bps
Output bytes :	0	0 bps
Input packets:	0	0 pps
Output packets:	0	0 pps

Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0

Flags: Trunk-Mode

show interfaces extensive

```
user@switch> show interfaces xe-0/0/1 extensive
```

```
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues       : 12 supported, 12 maximum usable queues
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped    : 2011-06-01 00:42:03 PDT (00:03:08 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
    Input bytes   : 0                0 bps
    Output bytes  : 0                0 bps
    Input packets : 0                0 pps
    Output packets: 0                0 pps
  IPv6 transit statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort   0                0                0
    1 fc7           0                0                0
    2 no-loss       0                0                0
    3 fcoe          0                0                0
    4 fc4           0                0                0
    5 fc5           0                0                0
    6 fc6           0                0                0
    7 network-cont  0                0                0
```

```

      8 mcast                                0                                0                                0
Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast

Active alarms   : None
Active defects  : None

MAC statistics:                                Receive                                Transmit
Total octets                                0                                0
Total packets                                0                                0
Unicast packets                             0                                0
Broadcast packets                           0                                0
Multicast packets                           0                                0
CRC/Align errors                             0                                0
FIFO errors                                 0                                0
MAC control frames                           0                                0
MAC pause frames                             0                                0
Oversized frames                             0
Jabber frames                               0
Fragment frames                             0
VLAN tagged frames                           0
Code violations                             0

MAC Priority Flow Control Statistics:
Priority : 0                                0                                0
Priority : 1                                0                                0
Priority : 2                                0                                0
Priority : 3                                0                                0
Priority : 4                                0                                0
Priority : 5                                0                                0
Priority : 6                                0                                0
Priority : 7                                0                                0

Filter statistics:
Input packet count                           0
Input packet rejects                         0
Input DA rejects                             0
Input SA rejects                             0
Output packet count                           0
Output packet pad count                       0

```

```

Output packet error count                                0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
    0 best-effort          75      7500000000    75      0      low
none
    7 network-control      5       500000000     5      0      low
none
    8 mcast                20      2000000000    20      0      low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input  bytes :                0
  Output bytes :                0
  Input  packets:                0
  Output packets:                0
Local statistics:
  Input  bytes :                0
  Output bytes :                0
  Input  packets:                0
  Output packets:                0
Transit statistics:
  Input  bytes :                0      0 bps
  Output bytes :                0      0 bps
  Input  packets:                0      0 pps
  Output packets:                0      0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces extensive (Asymmetric Flow Control)

user@switch> show interfaces xe-0/0/1 extensive

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:

```

```

Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags       : None
  CoS queues       : 12 supported, 12 maximum usable queues
  Hold-times       : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped    : 2011-06-01 00:42:03 PDT (00:03:08 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
  Traffic statistics:
    Input  bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input  packets:                0                0 pps
    Output packets:                0                0 pps
  IPv6 transit statistics:
    Input  bytes :                0
    Output bytes :                0
    Input  packets:                0
    Output packets:                0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 12 supported, 9 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets
    0 best-effort    0                0                0
    1 fc7            0                0                0
    2 no-loss        0                0                0
    3 fcoe           0                0                0
    4 fc4            0                0                0
    5 fc5            0                0                0
    6 fc6            0                0                0
    7 network-cont   0                0                0
    8 mcast          0                0                0
  Queue number:
    Mapped forwarding classes
    0 best-effort
    1 fc7
    2 no-loss
    3 fcoe
    4 fc4

```



```

5          fc5
6          fc6
7          network-control
8          mcast
Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive          Transmit
Total packets         0          0
Unicast packets       0          0
Broadcast packets     0          0
Multicast packets     0          0
CRC/Align errors      0          0
FIFO errors           0          0
MAC control frames    0          0
MAC pause frames      0          0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations        0
MAC Priority Flow Control Statistics:
Priority : 0           0          0
Priority : 1           0          0
Priority : 2           0          0
Priority : 3           0          0
Priority : 4           0          0
Priority : 5           0          0
Priority : 6           0          0
Priority : 7           0          0
Filter statistics:
Input packet count     0
Input packet rejects   0
Input DA rejects       0
Input SA rejects       0
Output packet count    0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue      Bandwidth      Buffer Priority  Limit

```

	%	bps	%	usec		
0 best-effort	75	7500000000	75	0	low	none
7 network-control	5	500000000	5	0	low	none
8 mcast	20	2000000000	20	0	low	none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Transit statistics:

Input bytes : 0 0 bps

Output bytes : 0 0 bps

Input packets: 0 0 pps

Output packets: 0 0 pps

Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0

Flags: Trunk-Mode

show interfaces terse

user@switch> show interfaces xe-0/0/1 terse

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/1	up	up			
xe-0/0/1.0	up	up	eth-switch		

show interfaces (QFabric System)

user@switch> show interfaces node1:xe-0/0/0

Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down

Interface index: 129, SNMP ifIndex: 2884086

Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU

Error: None, MAC-REWRITE Error: None,

Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```
Interface flags: Internal: 0x4000
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped    : Never
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
```

show interfaces statistics fabric

Syntax

```
show interfaces statistics fabric
<brief | detail | terse>
<descriptions>
<interface-name>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
```

Release Information

Command introduced in Junos OS Release 12.3 for the QFX Series.

Description

Display status information about the specified fabric interface.

Options

brief | detail | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

interface-name—(QFabric systems only) The interface name is either the serial number or the alias of the QFabric switch component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.

media—(Optional) Display media-specific information about network interfaces.

routing-instance (all | instance-name)—(Optional) Display all routing instances or the name of an individual routing instance.

snmp-index snmp-index—(Optional) Display information for the specified SNMP index of the interface.

Required Privilege Level

view

RELATED DOCUMENTATION

[Monitoring Interface Status and Traffic | 392](#)

[Troubleshooting Network Interfaces | 399](#)

[Troubleshooting an Aggregated Ethernet Interface](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

List of Sample Output

[show interfaces statistics fabric on page 1266](#)

[show interfaces statistics fabric brief on page 1276](#)

[show interfaces statistics fabric detail on page 1280](#)

[show interfaces statistics fabric terse on page 1282](#)

[show interfaces statistics fabric device-name on page 1283](#)

Output Fields

[Table 95 on page 1261](#) lists the output fields for the **show interfaces statistics fabric** command. Output fields are listed in the approximate order in which they appear.

Table 95: show interfaces statistics fabric Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP index number for the physical interface.	detail none
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Clocking	Reference clock source.	detail
Speed	Speed at which the interface is running.	All levels
Duplex	Duplex mode of the interface, either Full-Duplex or Half-Duplex.	All levels
MAC-REWRITE Error	Specifies if the encapsulation of the packet has been changed.	none
BPDU Error	Specifies if a BPDU has been received on a blocked interface.	none
Loopback	Loopback status: Enabled or Disabled. If loopback is enabled, type of loopback: Local or Remote.	All levels

Table 95: show interfaces statistics fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Source filtering	Source filtering status: Enabled or Disabled.	All levels
Flow control	Flow control status: Enabled or Disabled. This field is only displayed if asymmetric flow control is not configured.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
CoS queues	Number of CoS queues configured.	detail none
Hold-Times	Current interface hold-time up and hold-time down, in milliseconds.	detail
Current address	Configured MAC address.	detail none
Hardware address	Hardware MAC address.	detail none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) .	detail none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is Statistics last cleared: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, 2010-05-17 07:51:28 PDT (00:04:33 ago) .	detail
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail

Table 95: show interfaces statistics fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface:</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runt—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	detail none

Table 95: show interfaces statistics fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the fabric interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	detail none
Egress queues	Total number of egress queues supported on the specified interface.	detail
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified

Table 95: show interfaces statistics fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output rate	Output rate in bps and pps.	None specified
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail none
Flags	<p>Information about the logical interface.</p> <p>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.</p>	All levels
Input packets	Number of packets received on the interface.	detail none
Output packets	Number of packets transmitted on the interface.	detail none
Input packets	Number of packets received on the interface.	detail none
Output packets	Number of packets transmitted on the interface.	detail none
Encapsulation	<p>Encapsulation method used on the logical interface.</p> <ul style="list-style-type: none"> • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	All levels
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail

Table 95: show interfaces statistics fabric Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local statistics	<ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail
Transit statistics	<ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>NOTE: The bandwidth bps counter is not enabled.</p>	detail
Addresses, Flags	Information about the address flags.	detail none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
MTU	Maximum transmission unit size on the physical interface.	All levels
Destination	IP address of the remote side of the connection.	detail none
Local	IP address of the logical interface.	detail none
Broadcast	Broadcast address of the logical interlace.	detail none
Generation	Unique number for use by Juniper Networks technical support only.	detail
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail none

Sample Output

show interfaces statistics fabric

```
user@switch> show interfaces statistic fabric
```

Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
 Interface index: 49174, SNMP ifIndex: 1208484473
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:00, Hardware address: 00:00:00:00:00:00
 Last flapped : 2012-11-27 20:30:30 UTC (01:55:19 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/0.32768 (Index 83) (SNMP ifIndex 1208484474)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Down
 Interface index: 49175, SNMP ifIndex: 1208484363
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:04, Hardware address: 00:00:00:00:00:04
 Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/4.32768 (Index 84) (SNMP ifIndex 1208484364)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Down
 Interface index: 49176, SNMP ifIndex: 1208484367
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:06, Hardware address: 00:00:00:00:00:06
 Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/6.32768 (Index 85) (SNMP ifIndex 1208484368)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
 Interface index: 49177, SNMP ifIndex: 1208484479
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d
 Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1208484480)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Down
 Interface index: 49178, SNMP ifIndex: 1208484475
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:0f, Hardware address: 00:00:00:00:00:0f
 Last flapped : 2012-11-27 20:30:30 UTC (01:55:20 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-0/0/15.32768 (Index 87) (SNMP ifIndex 1208484476)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Down
 Interface index: 49211, SNMP ifIndex: 1208484377
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:02, Hardware address: 00:00:00:00:00:02
 Last flapped : 2012-11-27 20:30:47 UTC (01:55:03 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-1/0/2.32768 (Index 120) (SNMP ifIndex 1208484378)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Down
 Interface index: 49212, SNMP ifIndex: 1208484365
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPD
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:07, Hardware address: 00:00:00:00:00:07
 Last flapped : 2012-11-27 20:30:47 UTC (01:55:04 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-1/0/7.32768 (Index 121) (SNMP ifIndex 1208484366)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Down
 Interface index: 49213, SNMP ifIndex: 1208484625
 Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPD
 Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
 Disabled, Flow control: Disabled

Device flags : Present Running
 Interface flags: SNMP-Traps Internal: 0x0
 CoS queues : 12 supported, 12 maximum usable queues
 Current address: 00:00:00:00:00:0a, Hardware address: 00:00:00:00:00:0a
 Last flapped : 2012-11-27 20:30:47 UTC (01:55:04 ago)
 Statistics last cleared: Never
 Input rate : 0 bps (0 pps)
 Output rate : 0 bps (0 pps)
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:fte-1/0/10.32768 (Index 122) (SNMP ifIndex 1208484626)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
 Input packets : 0
 Output packets: 0
 Protocol eth-switch, MTU: 0

```

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up
  Interface index: 64, SNMP ifIndex: 1208483877
  Type: Ethernet, Link-level type: Ethernet, MTU: 1500
  Device flags   : Present Running
  Current address: 02:00:00:00:40:06, Hardware address: 02:00:00:00:40:01
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 26730
  Input errors: 0, Output errors: 0

Logical interface IC-WS001:bme0.0 (Index 4) (SNMP ifIndex 1208484065)
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  Input packets : 2715
  Output packets: 18
  Protocol inet, MTU: 1482
    Local: 128.0.32.0

Logical interface IC-WS001:bme0.1 (Index 5) (SNMP ifIndex 1208484091)
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 999
  Protocol inet, MTU: 1482
    Addresses
      Destination: 128/2, Local: 128.0.0.6, Broadcast: 191.255.255.255
      Destination: 128/2, Local: 128.0.130.2, Broadcast: 191.255.255.255

Logical interface IC-WS001:bme0.2 (Index 6) (SNMP ifIndex 1208484092)
  Flags: Encapsulation: ENET2
  Input packets : 180408
  Output packets: 23051
  Protocol inet, MTU: 1486
    Destination: 128/8, Local: 128.0.0.6, Broadcast: 128.255.255.255
    Destination: 128/8, Local: 128.0.130.2, Broadcast: 128.255.255.255
    Destination: 169.254/16, Local: 169.254.128.6, Broadcast: 169.254.255.255

    Destination: 169.254/16, Local: 169.254.193.1, Broadcast: 169.254.255.255

Physical interface: IC-WS001:bme1, Enabled, Physical link is Up
  Interface index: 49156, SNMP ifIndex: 1208483949
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex

```

Current address: 00:0d:0c:0f:00:03, Hardware address: 00:0d:0c:0f:00:03
 Last flapped : 1970-01-01 00:00:01 UTC (2238w5d 22:25 ago)
 Statistics last cleared: Never
 Input packets : 168885
 Output packets: 184712
 Input errors: 0, Output errors: 0

Logical interface IC-WS001:bmel.0 (Index 3) (SNMP ifIndex 1208483950)

Flags: Encapsulation: ENET2

Input packets : 168885

Output packets: 184712

Protocol inet, MTU: 1500

 Destination: 128/2, Local: 128.0.0.1, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.5, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.16, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.17, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.24, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.25, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.26, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.28, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.29, Broadcast: 191.255.255.255

 Destination: 128/2, Local: 128.0.0.31, Broadcast: 191.255.255.255

Protocol tnp, MTU: 1500

 Local: 0x5

Physical interface IC-WS001:dcfabric, Enabled, Physical link is Up

Interface index: 27, SNMP ifIndex: 1208484093

Type: Ethernet, Link-level type: Ethernet, MTU: 1572

Device flags : Present Running

Interface flags: SNMP-Traps

Current address: 00:0b:ca:fe:00:01, Hardware address: 00:0b:ca:fe:00:01

Last flapped : Never

Statistics last cleared: Never

 Input packets : 0

 Output packets: 0

Input errors: 0, Output errors: 0

Logical interface IC-WS001:dcfabric.0 (Index 64) (SNMP ifIndex 1208484094)

Flags: SNMP-Traps Encapsulation: ENET2

Input packets : 0

Output packets: 0

Protocol inet, MTU: 1558

Protocol mpls, MTU: 1546, Maximum labels: 3

Protocol eth-switch, MTU: 0

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up

Interface index: 66, SNMP ifIndex: 1208484104

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Link type : Full-Duplex

Current address: 00:23:9c:f1:a2:e6, Hardware address: 00:23:9c:f1:a2:e6

Last flapped : Never

Statistics last cleared: Never

Input packets : 1007238

Output packets: 63383

Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up

Interface index: 67, SNMP ifIndex: 1208484105

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Link type : Full-Duplex

Current address: 00:23:9c:f1:a2:e7, Hardware address: 00:23:9c:f1:a2:e7

Last flapped : Never

Statistics last cleared: Never

Input packets : 1007118

Output packets: 55381

Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down

Interface index: 68, SNMP ifIndex: 1208484106

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Link type : Full-Duplex

Current address: 00:23:9c:f1:a2:e8, Hardware address: 00:23:9c:f1:a2:e8

Last flapped : 2012-11-27 02:52:03 UTC (19:33:54 ago)

Statistics last cleared: Never

Input packets : 0

Output packets: 0

Input errors: 0, Output errors: 0

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down

Interface index: 69, SNMP ifIndex: 1208484107

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Speed: 1000mbps

Device flags : Present Running

```

Interface flags: SNMP-Traps
Link type       : Full-Duplex
Current address: 00:23:9c:f1:a2:e9, Hardware address: 00:23:9c:f1:a2:e9
Last flapped   : 2012-11-27 02:52:03 UTC (19:33:54 ago)
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

```

```

Physical interface: IC-WS001:vc0, Enabled, Physical link is Up
Interface index: 74, SNMP ifIndex: 1208484372
Type: Ethernet, Link-level type: 70, MTU: 1514, Speed: 1000mbps
Device flags   : Present Running
Link type      : Full-Duplex
Current address: 00:23:9c:f1:a2:e3, Hardware address: 00:23:9c:f1:a2:e3
Last flapped   : Never
Statistics last cleared: Never
  Input packets : 121842
  Output packets: 3548
Input errors: 0, Output errors: 0

```

```

Logical interface IC-WS001:vc0.32769 (Index 11) (SNMP ifIndex 1208484376)
Flags: LinkAddress 0-0 Encapsulation: ENET2
Input packets : 13044
Output packets: 3548

```

```

Physical interface: IC-WS001:vc1, Enabled, Physical link is Up
Interface index: 70, SNMP ifIndex: 1208484108
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
Device flags   : Present Running
Link type      : Full-Duplex
Current address: 00:23:9c:f1:a2:e6, Hardware address: 00:23:9c:f1:a2:e6
Last flapped   : Never
Statistics last cleared: Never
  Input packets : 767413
  Output packets: 46503
Input errors: 0, Output errors: 0

```

```

Logical interface IC-WS001:vc1.32768 (Index 7) (SNMP ifIndex 1208484109)
Flags: LinkAddress 0-0 Encapsulation: ENET2
Input packets : 735889
Output packets: 46503

```

```

Physical interface: IC-WS001:vc2, Enabled, Physical link is Up

```

```

Interface index: 71, SNMP ifIndex: 1208484369
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
Device flags    : Present Running
Link type       : Full-Duplex
Current address: 00:23:9c:f1:a2:e7, Hardware address: 00:23:9c:f1:a2:e7
Last flapped    : Never
Statistics last cleared: Never
  Input packets : 831710
  Output packets: 44548
Input errors: 0, Output errors: 0

Logical interface IC-WS001:vcp2.32768 (Index 8) (SNMP ifIndex 1208484373)
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  Input packets : 737844
  Output packets: 44548

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
Interface index: 72, SNMP ifIndex: 1208484370
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
Device flags    : Present Running
Link type       : Full-Duplex
Current address: 00:23:9c:f1:a2:e8, Hardware address: 00:23:9c:f1:a2:e8
Last flapped    : 2012-11-27 20:31:36 UTC (01:54:21 ago)
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

Logical interface IC-WS001:vcp3.32768 (Index 9) (SNMP ifIndex 1208484374)
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
Interface index: 73, SNMP ifIndex: 1208484371
Type: Ethernet, Link-level type: 70, MTU: 1492, Speed: 1000mbps
Device flags    : Present Running
Link type       : Full-Duplex
Current address: 00:23:9c:f1:a2:e9, Hardware address: 00:23:9c:f1:a2:e9
Last flapped    : 2012-11-27 20:31:36 UTC (01:54:21 ago)
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

```

```

Logical interface IC-WS001:vcp4.32768 (Index 10) (SNMP ifIndex 1208484375)
  Flags: LinkAddress 0-0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0

```

show interfaces statistics fabric brief

```
user@switch> show interfaces statistics fabric brief
```

```

Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

```

```

Logical interface IC-WS001:fte-0/0/0.32768
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

```

```

Physical interface: IC-WS001:fte-0/0/4, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

```

```

Logical interface IC-WS001:fte-0/0/4.32768
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

```

```

Physical interface: IC-WS001:fte-0/0/6, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running
  Interface flags: SNMP-Traps Internal: 0x0

```

```

Logical interface IC-WS001:fte-0/0/6.32768
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

```

```

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Device flags    : Present Running

```

Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/13.32768

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-0/0/15, Enabled, Physical link is Down

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-0/0/15.32768

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-1/0/2, Enabled, Physical link is Down

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/2.32768

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-1/0/7, Enabled, Physical link is Down

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/7.32768

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
eth-switch

Physical interface: IC-WS001:fte-1/0/10, Enabled, Physical link is Down

Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex,
Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0

Logical interface IC-WS001:fte-1/0/10.32768

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

eth-switch

Physical interface: IC-WS001:bme0, Enabled, Physical link is Up
 Type: Ethernet, Link-level type: Ethernet, MTU: 1500, Clocking: Unspecified,
 Speed: Unspecified

Device flags : Present Running

Logical interface IC-WS001:bme0.0

Flags: LinkAddress 0-0 Encapsulation: ENET2
 inet 128.0.32.0 --> 0/0

Logical interface IC-WS001:bme0.1

Flags: LinkAddress 0-0 Encapsulation: ENET2
 inet 128.0.0.6/2
 128.0.130.2/2

Logical interface IC-WS001:bme0.2

Flags: Encapsulation: ENET2
 inet 128.0.0.6/8
 128.0.130.2/8
 169.254.128.6/16
 169.254.193.1/16

Physical interface: IC-WS001:bme1, 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 IC-WS001:bme1.0

Flags: Encapsulation: ENET2
 inet 128.0.0.1/2
 128.0.0.5/2
 128.0.0.16/2
 128.0.0.17/2
 128.0.0.24/2
 128.0.0.25/2
 128.0.0.26/2
 128.0.0.28/2
 128.0.0.29/2
 128.0.0.31/2

tnp 0x5

Physical interface: IC-WS001:dcfabric, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1572, Clocking: Unspecified,
Speed: Unspecified

Device flags : Present Running

Interface flags: SNMP-Traps

Logical interface IC-WS001:dcfabric.0

Flags: SNMP-Traps Encapsulation: ENET2

inet

mpls

eth-switch

Physical interface: IC-WS001:pme0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme1, Enabled, Physical link is Up

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme2, Enabled, Physical link is Down

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:pme3, Enabled, Physical link is Down

Type: Ethernet, Link-level type: Ethernet, MTU: 1510, Clocking: Unspecified,
Speed: 1000mbps

Device flags : Present Running

Interface flags: SNMP-Traps

Physical interface: IC-WS001:vcp0, Enabled, Physical link is Up

Type: Ethernet, Link-level type: 70, MTU: 1514, Clocking: Unspecified, Speed:
1000mbps

Device flags : Present Running

Logical interface IC-WS001:vcp0.32769

Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcpl, Enabled, Physical link is Up

```

    Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
    Device flags      : Present Running

    Logical interface IC-WS001:vcpl.32768
        Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcpl2, Enabled, Physical link is Up
    Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
    Device flags      : Present Running

    Logical interface IC-WS001:vcp2.32768
        Flags: LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp3, Enabled, Physical link is Down
    Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
    Device flags      : Present Running

    Logical interface IC-WS001:vcp3.32768
        Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

Physical interface: IC-WS001:vcp4, Enabled, Physical link is Down
    Type: Ethernet, Link-level type: 70, MTU: 1492, Clocking: Unspecified, Speed:
1000mbps
    Device flags      : Present Running

    Logical interface IC-WS001:vcp4.32768
        Flags: Device-Down LinkAddress 0-0 Encapsulation: ENET2

```

show interfaces statistics fabric detail

user@switch> **show interfaces statistics fabric detail**

```

show interfaces statistics fabric detail
Physical interface: IC-WS001:fte-0/0/0, Enabled, Physical link is Down
    Interface index: 49174, SNMP ifIndex: 1208484473, Generation: 153
    Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Disabled
    Device flags      : Present Running
    Interface flags: SNMP-Traps Internal: 0x0
    CoS queues        : 12 supported, 12 maximum usable queues

```



```

Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:00:00:00:00:00, Hardware address: 00:00:00:00:00:00
Last flapped   : 2012-11-27 20:30:30 UTC (02:04:59 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes   : 0 0 bps
  Output bytes  : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 5 in use
Queue counters:
  Queued packets  Transmitted packets  Dropped packets
0 fabric_fcset   0 0 0
1 fabric_fcset   0 0 0
2 fabric_fcset   0 0 0
3 fabric_fcset   0 0 0
4 fabric_fcset   0 0 0
5 fabric_fcset   0 0 0
6 fabric_fcset   0 0 0
7 fabric_fcset   0 0 0
8 fabric_fcset   0 0 0
9 fabric_fcset   0 0 0
10 fabric_fcset  0 0 0
11 fabric_fcset  0 0 0

Logical interface IC-WS001:fte-0/0/0.32768 (Index 83) (SNMP ifIndex 1208484474)
(Generation 148)
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
Protocol eth-switch, MTU: 0, Generation: 176, Route table: 0

```

show interfaces statistics fabric terse

user@switch> show interfaces statistics fabric terse

Interface	Admin	Link	Proto	Local	Remote
IC-WS001:fte-0/0/0	up	down			
IC-WS001:fte-0/0/0.32768	up	down	eth-switch		
IC-WS001:fte-0/0/4	up	down			
IC-WS001:fte-0/0/4.32768	up	down	eth-switch		
IC-WS001:fte-0/0/6	up	down			
IC-WS001:fte-0/0/6.32768	up	down	eth-switch		
IC-WS001:fte-0/0/13	up	down			
IC-WS001:fte-0/0/13.32768	up	down	eth-switch		
IC-WS001:fte-0/0/15	up	down			
IC-WS001:fte-0/0/15.32768	up	down	eth-switch		
IC-WS001:fte-1/0/2	up	down			
IC-WS001:fte-1/0/2.32768	up	down	eth-switch		
IC-WS001:fte-1/0/7	up	down			
IC-WS001:fte-1/0/7.32768	up	down	eth-switch		
IC-WS001:fte-1/0/10	up	down			
IC-WS001:fte-1/0/10.32768	up	down	eth-switch		
IC-WS001:bme0	up	up			
IC-WS001:bme0.0	up	up	inet	128.0.32.0	--> 0/0
IC-WS001:bme0.1	up	up	inet	128.0.0.6/2	
				128.0.130.2/2	
IC-WS001:bme0.2	up	up	inet	128.0.0.6/8	
				128.0.130.2/8	
				169.254.128.6/16	
				169.254.193.1/16	
IC-WS001:bme1	up	up			

```

IC-WS001:bme1.0      up    up    inet    128.0.0.1/2
                   128.0.0.5/2
                   128.0.0.16/2
                   128.0.0.17/2
                   128.0.0.24/2
                   128.0.0.25/2
                   128.0.0.26/2
                   128.0.0.28/2
                   128.0.0.29/2
                   128.0.0.31/2
                   tnp    0x5
IC-WS001:dcfabric    up    up
IC-WS001:dcfabric.0  up    up    inet
                   mpls
                   eth-switch
IC-WS001:pme0        up    up
IC-WS001:pme1        up    up
IC-WS001:pme2        up    down
IC-WS001:pme3        up    down
IC-WS001:vcp0        up    up
IC-WS001:vcp0.32769  up    up
IC-WS001:vcp1        up    up
IC-WS001:vcp1.32768  up    up
IC-WS001:vcp2        up    up
IC-WS001:vcp2.32768  up    up
IC-WS001:vcp3        up    down
IC-WS001:vcp3.32768  up    down
IC-WS001:vcp4        up    down
IC-WS001:vcp4.32768  up    down

```

show interfaces statistics fabric device-name

user@switch> show interfaces statistics fabric IC-WS001:fte-0/0/13

```

Physical interface: IC-WS001:fte-0/0/13, Enabled, Physical link is Down
  Interface index: 49177, SNMP ifIndex: 1208484479
  Link-level type: Ethernet, MTU: 9232, Speed: 40Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Disabled
  Device flags      : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  CoS queues        : 12 supported, 12 maximum usable queues
  Current address: 00:00:00:00:00:0d, Hardware address: 00:00:00:00:00:0d
  Last flapped      : 2012-11-27 20:30:30 UTC (02:09:53 ago)

```

```
Statistics last cleared: Never
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
Input errors: 0, Output errors: 0
```

```
Logical interface IC-WS001:fte-0/0/13.32768 (Index 86) (SNMP ifIndex 1208484480)
```

```
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Input packets : 0
Output packets: 0
Protocol eth-switch, MTU: 0
```

show interfaces vlan

Syntax

```
show interfaces (vlan | vlan.vlan-id)
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | instance-name)>
<snmp-index snmp-index>
<statistics>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display status information about routed VLAN interfaces (RVIs).

Options

vlan | vlan.vlan-id—Display status information for the specified RVI.

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.

routing-instance (all | instance-name)—(Optional) Associate this RVI with the named routing instance.

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

show ethernet-switching table

show vlans

Monitoring Interface Status and Traffic

Troubleshooting Network Interfaces on EX3200 Switches

Troubleshooting Network Interfaces on EX4200 Switches

List of Sample Output[show interfaces vlan on page 1296](#)[show interfaces vlan terse on page 1296](#)[show interfaces vlan extensive on page 1297](#)[show interfaces vlan detail on page 1298](#)**Output Fields**

Table 96 on page 1286 lists the output fields for the **show interfaces vlan** command. Output fields are listed in the approximate order in which they appear. The level of output *none* means the basic command with no optional options—that is, either just **show interfaces vlan** or **show interfaces vlan.vlan-id**.

Table 96: show interfaces vlan Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, which is always vlan .	All levels
Enabled	State of the interface: Enabled or Disabled , followed by the statement Physical link is <Up/Down>	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Because this is routed VLAN interface information, this entry is always VLAN .	detail extensive none
Link-level type	Encapsulation (added control information) being used on the physical interface. Because this is routed VLAN interface information, this entry is always VLAN .	All levels

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
MTU	Maximum transmission unit (MTU) size on the physical interface. The default MTU size depends on the switch platform. Changing either the media MTU or protocol MTU causes an interface to be deleted and added again.	All levels
Clocking	Value is always Unspecified —not applicable on switches.	detail extensive
Speed	Speed of the interface, either Auto if autonegotiation of speed is enabled or a number representing the configured speed in megabits per second.	detail extensive none
Device flags	<p>Information about the physical device such as:</p> <p>Dest-route-down—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status.</p> <p>Down—Device has been administratively disabled.</p> <p>Hear-Own-Xmit—Device receives its own transmissions.</p> <p>Is-Default—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities.</p> <p>Is-Preferred—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet.</p> <p>Is-Primary—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface.</p> <p>Link-Layer-Down—The link-layer protocol has failed to connect with the remote endpoint.</p> <p>Loopback—Switch is in physical loopback.</p> <p>Loop-Detected—The link layer has received frames that it sent, thereby detecting a physical loopback.</p> <p>No-Carrier—On media that support carrier recognition, no carrier is currently detected.</p> <p>No-Multicast—Device does not support multicast traffic.</p> <p>Preferred—This address is a candidate to become the preferred address.</p> <p>Present—Device is physically present and recognized.</p> <p>Promiscuous—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</p> <p>Primary—This address is a candidate to become the primary address.</p> <p>Quench—Transmission on the device is quenched, because the output buffer is overflowing</p> <p>Recv-All-Multicasts—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</p> <p>Running—Device is active and enabled.</p>	detail extensive none

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link type	Link mode of the interface— Auto if autonegotiation is enabled, or the configured Full-Duplex or Half-Duplex .	detail extensive none
Link flags	Value is always None —not applicable on switches.	detail extensive none
Physical Info	Value is always Unspecified —not applicable on switches.	detail extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	MAC address of the hardware.	detail extensive none
Hardware address	MAC address of the switch.	detail extensive none
Alternate link address	Value is always Unspecified —not applicable on switches.	detail extensive
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) . The entry can also be Never .	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive none

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets transmitted or received on the physical interface for supported switches.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter on EX3200 switches and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <i>l3-interface-ingress-counting</i>. • Output bytes—Number of bytes sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • Input packets—Number of packets received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <i>l3-interface-ingress-counting</i>. • Output packets—Number of packets sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. 	detail extensive
IPv6 transit statistics	<p>Number and rate of bytes and packets transmitted and/or received on the IPv6 interface for supported switches.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <i>l3-interface-ingress-counting</i>. • Output bytes—Number of bytes sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • Input packets—Number of packets received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command <i>l3-interface-ingress-counting</i>. • Output packets—Number of packets sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for and EX8200 switches. 	detail extensive

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input Errors	<p>Input errors on the interface. The following paragraphs explain some of the counters whose meaning may not be obvious.</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this value increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC. If this value is ever nonzero, the interface is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This value 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 interface 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. Both Gigabit Ethernet interfaces and 10 Gigabit Ethernet interfaces support only full-duplex operation, so for those two interfaces, this value should always be zero. If the value is nonzero for either Gigabit Ethernet or 10 Gigabit Ethernet, 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 interface. If this value is ever nonzero, the interface is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
vlan.vlan-id, Index, SNMP ifIndex	VLAN ID, index, and SNMP index number for the logical interface. The logical interface index values reflect the item's initialization sequence.	detail extensive none
Generation	Unique number for Juniper Networks Technical support use only.	detail extensive none

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	<p>Errors that have occurred on this interface, such as Link Layer Down. Other possible flags include:</p> <ul style="list-style-type: none"> • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Hardware-Down—Interface protocol initialization failed to complete successfully. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational. 	detail extensive none
SNMP-Traps	Each configured SNMP trap has a number that appears here—0x0 is always displayed for logical interface SNMP traps.	detail extensive none
Encapsulation	Encapsulation method, which is the process of adding control information. The value is always Ethernet 2 (ENET2) for logical encapsulation.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the logical interface of supported switches. Traffic statistics represent the sum of the next two fields, Local statistics and Transit statistics. Note that these are not the values for the RVI ingress or egress counters—for that value, see Transit statistics below.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. Same value as the physical interface. • Output bytes—Number of bytes sent on the interface. Same value as the physical interface. • Input packets—Number of packets received on the interface. Same value as the physical interface. • Output packets—Number of bytes sent on the interface. Same value as the physical interface. <p>NOTE: The bandwidth bps counter is not enabled on the switches.</p>	detail extensive

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local statistics	<p>Number and rate of bytes and packets received and transmitted locally by the Routing Engine on the logical interface of supported switches. All packets for protocols and process statistics are counted here.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. Same value as for the physical interface. • Output bytes—Number of bytes sent on the interface. Same value as for the physical interface. • Input packets—Number of packets received on the interface. Same value as for the physical interface. • Output packets—Number of bytes sent on the interface. Same value as for the physical interface. 	detail extensive none
Transit statistics	<p>Number and rate of bytes and packets received and transmitted on the RVI logical interface of supported switches. Look at this value to see the RVI ingress and egress count.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • Output bytes—Number of bytes sent on the interface. This egress counter is automatic for EX8200. • Input packets—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • Output packets—Number of packets sent on the interface. This egress counter is automatic for EX8200 switches. 	detail extensive

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 transit statistics	<p>Number and rate of IPv6 bytes and packets received and transmitted on the RVI logical interface of supported switches. Transit values are unique to the logical interface and do not appear in physical interface output. Look at the values listed below to see the RVI ingress and egress count for IPv6 traffic.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • Output bytes—Number of bytes sent by the interface. This egress counter is automatic for EX8200 switches. • Input packets—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 and configurable for EX8200 switches. • Output packets—Number of packets sent by the interface. This egress counter is automatic for EX8200 switches. <p>NOTE: The bandwidth bps counter is not enabled on the switches.</p>	detail extensive
Protocol	Protocol used for the logical interface—this value is inet for IPv4 traffic and inet6 for IPv6 traffic.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 .	detail extensive none
Protocol flags	Information about the protocol such as Targeted-broadcast .	detail extensive none

Table 96: show interfaces vlan Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protocol addresses and Address flags	<p>Protocol address values here can be:</p> <p>Dest-route-down—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status</p> <p>Device-down—Device has been administratively disabled.</p> <p>Disabled—Interface is administratively disabled.</p> <p>Down—A hardware failure has occurred.</p> <p>Hardware-Down—Interface protocol initialization failed to complete successfully.</p> <p>Is-Default—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities.</p> <p>Is-Preferred—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet.</p> <p>Is-Primary—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface.</p> <p>Preferred—This address is a candidate to become the preferred address.</p> <p>Primary—This address is a candidate to become the primary address.</p> <p>SNMP-Traps—SNMP trap notifications are enabled.</p> <p>Up—Interface is enabled and operational.</p>	<p>detail</p> <p>extensive</p> <p>none</p>
Address destination	Logical destination's network address.	<p>detail</p> <p>extensive</p> <p>none</p>
Local address	IP address of the logical interface.	<p>detail</p> <p>extensive</p> <p>none</p>
Broadcast address	Broadcast address of the logical interface.	<p>detail</p> <p>extensive</p> <p>none</p>
Generation	Unique number for use by Juniper Networks technical support only.	<p>detail</p> <p>extensive</p>

Sample Output

show interfaces vlan

user@switch> show interfaces vlan

```
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Speed: 1000mbps
  Device flags      : Present Running
  Link type         : Full-Duplex
  Link flags        : None
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  Last flapped     : Never
    Input packets : 0
    Output packets: 0

Logical interface vlan.0 (Index 82) (SNMP ifIndex 557)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 1
  Protocol inet
    Flags: Targeted-broadcast
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 1
  Protocol inet
    Flags: Targeted-broadcast
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255
```

show interfaces vlan terse

user@switch> show interfaces vlan terse

Interface	Admin	Link	Proto	Local	Remote
vlan	up	up			
vlan.0	up	down	inet	10.1.1.1/24	
vlan.1	up	down	inet	10.1.2.1/24	

show interfaces vlan extensive

```
user@switch> show interfaces vlan extensive
```

```
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556, Generation: 153
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags      : Present Running
  Link type         : Full-Duplex
  Link flags        : None
  Physical info     : Unspecified
  Hold-times        : Up 0 ms, Down 0 ms
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  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 vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes   :                0
    Output bytes   :               42
    Input  packets :                0
    Output packets :                1
  Local statistics:
    Input  bytes   :                0
    Output bytes   :               42
    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
Protocol inet, Generation: 159, Route table: 0
  Flags: Targeted-broadcast
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
    Generation: 138

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558) (Generation 148)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes   :                0
    Output bytes   :               42
    Input  packets :                0
    Output packets :                1
  Local statistics:
    Input  bytes   :                0
    Output bytes   :               42
    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
Protocol inet, Generation: 160, Route table: 0
  Flags: Targeted-broadcast
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255,
    Generation: 140

```

show interfaces vlan detail

user@switch> **show interfaces vlan detail**

```

Physical interface: vlan, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 556, Generation: 153
Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
Speed: 1000mbps

```

```

Device flags      : Present Running
Link type        : Full-Duplex
Link flags       : None
Physical info    : Unspecified
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
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

Logical interface vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input  bytes :                0
    Output bytes :               42
    Input  packets:              0
    Output packets:              1
  Local statistics:
    Input  bytes :                0
    Output bytes :               42
    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
  Protocol inet, Generation: 159, Route table: 0
    Flags: Targeted-broadcast
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
      Generation: 138

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558) (Generation 148)

```

```
Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input  bytes   :                0
  Output bytes   :               42
  Input  packets :                0
  Output packets :                1
Local statistics:
  Input  bytes   :                0
  Output bytes   :               42
  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
Protocol inet, Generation: 160, Route table: 0
  Flags: Targeted-broadcast
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255,
    Generation: 140
```

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

<i>Configuring Aggregated Ethernet Links (CLI Procedure)</i>
Configuring Link Aggregation 281
Configuring Aggregated Ethernet LACP (CLI Procedure) 283
Configuring Aggregated Ethernet LACP (CLI Procedure) 283
Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches 313
Understanding Aggregated Ethernet Interfaces and LACP for Switches 262
<i>Junos OS Interfaces Fundamentals Configuration Guide</i>

List of Sample Output

[show lacp interfaces \(Aggregated Ethernet\) on page 1306](#)

[show lacp interfaces \(Gigabit Ethernet\) on page 1307](#)

[show lacp interfaces \(10 Gigabit Ethernet\) on page 1307](#)

Output Fields

[Table 97 on page 1302](#) lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 97: 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 97: 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"> • Role—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> • Actor—Local device participating in LACP negotiation. • Partner—Remote device participating in LACP negotiation. • Exp—Expired state. Yes indicates the actor or partner is in an expired state. No indicates the actor or partner is not in an expired state. • Def—Default. Yes indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. No indicates the operational partner information in use has been received in an LACP PDU. • Dist—Distribution of outgoing frames. No indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is Yes. • Col—Collection of incoming frames. Yes indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is No. • Syn—Synchronization. If the value is Yes, the link is considered 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 No, the link is not synchronized. It is currently not in the right aggregation. • Aggr—Ability of aggregation port to aggregate (Yes) or to operate only as an individual link (No). • Timeout—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (Slow Timeout or Fast Timeout). In a fast timeout, PDUs are sent every second and in a slow timeout, PDUs are sent every 30 seconds. LACP timeout occurs when 3 consecutive PDUs are missed. If LACP timeout is a fast timeout, the time taken when 3 consecutive PDUs are missed is 3 seconds (3x1 second). If LACP timeout is a slow timeout, the time taken is 90 seconds (3x30 seconds). • Activity—Actor or partner's port activity. Passive indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is Active. Active indicates the port's preference to participate in the protocol regardless of the partner's control value. • Core isolation state down (CDN)— LACP interface state. Down indicates the LACP interface is down because all the eBGP sessions for Ethernet VPN (EVPN) are down. 	All Levels

Table 97: show lacp interfaces Output Fields (continued)

Field Name	Field Description	Level of Out
LACP Protocol		All Levels

Table 97: 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"> • Link state (active or standby) indicated in parentheses next to the interface when link protection is configured. • Receive State—One of the following values: <ul style="list-style-type: none"> • Current—The state machine receives an LACP PDU and enters the Current state. • Defaulted—If no LACP PDU is received before the timer for the Current state expires a second time, the state machine enters the Defaulted state. • Expired—If no LACP PDU is received before the timer for the Current state expires once, the state machine enters the Expired state. • Initialize—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the Initialize state. • LACP Disabled—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to LACP Disabled. This state is similar to the Defaulted state, except that the port is forced to operate as an individual port. • Port Disabled—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the Port Disabled state. • Transmit State—Transmit state of state machine. One of the following values: <ul style="list-style-type: none"> • Fast Periodic—Periodic transmissions are enabled at a fast transmission rate. • No Periodic—Periodic transmissions are disabled. • Periodic Timer—Transitory state entered when the periodic timer expires. • Slow Periodic—Periodic transmissions are enabled at a slow transmission rate. • Mux State—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> • Attached—Multiplexer state machine initiates the process of attaching the port to the selected aggregator. • Collecting—Yes indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. No indicates the receive function of this link is not enabled. • Collecting Distributing—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution. • Detached—Process of detaching the port from the aggregator is in progress. • Distributing—Yes indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. No indicates the transmit function of this link is not enabled. 	

Table 97: show lacp interfaces Output Fields (continued)

Field Name	Field Description	Level of Out
	<ul style="list-style-type: none"> ● Waiting—Multiplexer state machine is in a holding process, awaiting an outcome. 	
LACP info	<ul style="list-style-type: none"> ● Role can be one of the following: <ul style="list-style-type: none"> ● Actor—Local device participating in LACP negotiation. ● Partner—Remote device participating in LACP negotiation. ● System priority—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer. ● System identifier—Actor or partner system ID, encoded as a MAC address. ● Port priority—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer. ● Port number—Port number assigned to the port by the actor or partner, encoded as an unsigned integer. ● Port key—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer. 	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									
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 lacp statistics interfaces (View)

Syntax

```
show lacp statistics interfaces interface-name
```

Release Information

Command modified in Release 10.2 of Junos OS.
Command introduced in Release 11.1 of Junos OS for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.

Options

interface-name—(Optional) Name of an interface.

Required Privilege Level

view

RELATED DOCUMENTATION

Verifying LACP on Redundant Ethernet Interfaces
Verifying the Status of a LAG Interface
Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets 320
Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch 299
Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch 328

List of Sample Output

[show lacp statistics interfaces on page 1309](#)

Output Fields

[Table 98 on page 1309](#) lists the output fields for the **show lacp statistics interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 98: show lacp statistics interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.
LACP Statistics	<p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> • LACP Rx—counter that increments for each received LACP packet. • LACP Tx—counter that increments for each transmitted LACP packet. • Unknown Rx—number of unrecognized packet errors logged. • Illegal Rx—number of invalid packets received. <p>NOTE: Starting in Junos OS Evolved Release 18.3R1, the clear interfaces statistics command clears LACP statistics as well as the counters displayed in the show lacp statistics interfaces command.</p>

Sample Output

show lacp statistics interfaces

user@host> **show lacp statistics interfaces ae0**

Aggregated interface: ae0

LACP Statistics:	LACP Rx	LACP Tx	Unknown Rx	Illegal Rx
ge-2/0/0	1352	2035	0	0
ge-2/0/1	1352	2056	0	0
ge-2/2/0	1352	2045	0	0
ge-2/2/1	1352	2043	0	0

show redundant-trunk-group

Syntax

```
show redundant-trunk-group <group-name group-name>
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.

Description

Display information about redundant trunk groups.

Options

group-name group-name—Display information about the specified redundant trunk group.

Required Privilege Level

view

RELATED DOCUMENTATION

- [Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support](#)
- [Understanding Redundant Trunk Links \(Legacy RTG Configuration\)](#)

List of Sample Output

[show redundant-trunk-group group-name Group1 on page 1311](#)

Output Fields

[Table 99 on page 1310](#) lists the output fields for the **show redundant-trunk-group** command. Output fields are listed in the approximate order in which they appear.

Table 99: show redundant-trunk-group Output Fields

Field Name	Field Description
Group name	Name of the redundant trunk port group.
Interface	Name of an interface belonging to the trunk port group.

Table 99: show redundant-trunk-group Output Fields (*continued*)

Field Name	Field Description
State	<p>Operating state of the interface.</p> <ul style="list-style-type: none"> • Up denotes the interface is up. • Down denotes the interface is down. • Pri denotes a primary interface. • Act denotes an active interface.
Time of last flap	Date and time at which the advertised link became unavailable, and then, available again.
Flap count	Total number of flaps since the last switch reboot.

Sample Output

```
show redundant-trunk-group group-name Group1
```

```
user@switch> show redundant-trunk-group group-name Group1
```

Group name	Interface	State	Time of last flap	Flap Count
Group1	ge-0/0/45.0	UP/Pri/Act	Never	0
	ge-0/0/47.0	UP	Never	0

show uplink-failure-detection

Syntax

```
show uplink-failure-detection
<group group-name>
```

Release Information

Command introduced in Junos OS Release 11.1 for EX Series switches.
 Command introduced in Junos OS Release 12.1 for QFX Series switches.

Description

Display information about the uplink-failure-detection group, the member interfaces, and their status.

Options

- none**—Display information about all groups configured for uplink failure detection.
- group group-name**—(Optional) Display information about the specified group only.

Required Privilege Level

view

RELATED DOCUMENTATION

Overview of Uplink Failure Detection 215
Configuring Interfaces for Uplink Failure Detection 217
Example: Configuring Interfaces for Uplink Failure Detection 218

List of Sample Output

- [show uplink-failure-detection on page 1313](#)
- [show uplink-failure-detection group g2 on page 1314](#)

Output Fields

[Table 100 on page 1312](#) lists the output fields for the **show uplink-failure-detection** command. Output fields are listed in the approximate order in which they appear.

Table 100: show uplink-failure-detection Output Fields

Field Name	Field Description
Group	Name of the group.

Table 100: show uplink-failure-detection Output Fields (*continued*)

Field Name	Field Description
Uplink	The uplink interface or interfaces configured as link-to-monitor. NOTE: The asterisk (*) indicates that the link is up.
Downlink	The downlink interface or interfaces configured as link-to-disable. NOTE: The asterisk (*) indicates that the link is up.
Failure Action	Status of uplink failure detection: <ul style="list-style-type: none"> • Active—The switch has detected an uplink failure and has brought the downlink down. • Inactive—The uplink or uplinks are up.
Debounce Interval	The amount of time, in seconds, that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces.

Sample Output

show uplink-failure-detection

```
user@switch> show uplink-failure-detection
```

```

Group          : group1
Uplink         : ge-0/0/0*
Downlink       : ge-0/0/1*
Failure Action  : Inactive

Group          : group2
Uplink         : ge-0/0/3.0
Downlink       : ge-0/0/4.0
Failure Action  : Active

Debounce Interval : 20
```

Sample Output

show uplink-failure-detection group g2

user@switch> **show uplink-failure-detection group g2**

```
Group           : group1
Uplink           : ge-0/0/0*
Downlink         : ge-0/0/1*
Failure Action   : Inactive

Group           : group2
Uplink           : ge-0/0/3.0
Downlink         : ge-0/0/4.0
Failure Action   : Active

Debounce Interval : 20
```

show virtual-chassis vc-port diagnostics optics

Syntax

```
show virtual-chassis vc-port diagnostics optics  
<all-members | local | member member-id>  
<interface-name>
```

Release Information

Command introduced in Junos OS Release 12.2 for EX Series switches.

Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description

Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.

On some EX Series switches, the **request virtual-chassis vc-port diagnostics optics** command must be entered to run a diagnostic scan before you can gather the **show virtual-chassis vc-port diagnostics optics** output.

Options

none—Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.

all-members—(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.

interface-name—(Optional) Display diagnostics information for the transceiver installed in a specified VCP.

local—(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine where you enter this command.

member *member-id*—(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.

Required Privilege Level

view

RELATED DOCUMENTATION

show virtual-chassis vc-port

Install a Transceiver

Remove a Transceiver

Junos OS Ethernet Interfaces Configuration Guide

List of Sample Output

[show virtual-chassis vc-port diagnostics optics on page 1319](#)

[show virtual-chassis vc-port diagnostics optics \(interface-name\) on page 1326](#)

[show virtual-chassis vc-port diagnostics optics local on page 1329](#)

[show virtual-chassis vc-port diagnostics optics \(member member-id\) on page 1331](#)

Output Fields

[Table 101 on page 1316](#) lists the output fields for the **show virtual-chassis vc-port diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

Table 101: show virtual-chassis vc-port diagnostics optics Output Fields

Field Name	Field Description
FPC	Displays the FPC slot number.
Virtual chassis port	Displays the name of the VCP.
Laser bias current	Displays the 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	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is <i>On</i> or <i>Off</i> .
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is <i>On</i> or <i>Off</i> .
Laser bias current high warning	Displays whether the laser bias power setting high warning is <i>On</i> or <i>Off</i> .

Table 101: show virtual-chassis vc-port diagnostics optics Output Fields (continued)

Field Name	Field Description
Laser bias current low warning	Displays whether the laser bias power setting low warning is <i>On</i> or <i>Off</i> .
Laser output power high alarm	Displays whether the laser output power high alarm is <i>On</i> or <i>Off</i> .
Laser output power low alarm	Displays whether the laser output power low alarm is <i>On</i> or <i>Off</i> .
Laser output power high warning	Displays whether the laser output power high warning is <i>On</i> or <i>Off</i> .
Laser output power low warning	Displays whether the laser output power low warning is <i>On</i> or <i>Off</i> .
Module temperature high alarm	Displays whether the module temperature high alarm is <i>On</i> or <i>Off</i> .
Module temperature low alarm	Displays whether the module temperature low alarm is <i>On</i> or <i>Off</i> .
Module temperature high warning	Displays whether the module temperature high warning is <i>On</i> or <i>Off</i> .
Module temperature low warning	Displays whether the module temperature low warning is <i>On</i> or <i>Off</i> .
Module voltage high alarm	Displays whether the module voltage high alarm is <i>On</i> or <i>Off</i> .
Module voltage low alarm	Displays whether the module voltage low alarm is <i>On</i> or <i>Off</i> .
Module voltage high warning	Displays whether the module voltage high warning is <i>On</i> or <i>Off</i> .
Module voltage low warning	Displays whether the module voltage low warning is <i>On</i> or <i>Off</i> .
Laser rx power high alarm	Displays whether the receive laser power high alarm is <i>On</i> or <i>Off</i> .
Laser rx power low alarm	Displays whether the receive laser power low alarm is <i>On</i> or <i>Off</i> .
Laser rx power high warning	Displays whether the receive laser power high warning is <i>On</i> or <i>Off</i> .
Laser rx power low warning	Displays whether the receive laser power low warning is <i>On</i> or <i>Off</i> .
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.

Table 101: show virtual-chassis vc-port diagnostics optics Output Fields (continued)

Field Name	Field Description
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.

Table 101: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show virtual-chassis vc-port diagnostics optics

user@switch> **show virtual-chassis vc-port diagnostics optics**

```
fpc0:
-----
Virtual chassis port: vcp-0
    Optical diagnostics                :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                :  N/A

fpc1:
-----
Virtual chassis port: vcp-0
    Optical diagnostics                :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                :  N/A

fpc2:
-----
Virtual chassis port: vcp-2/0
    Optical diagnostics                :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                :  N/A
Virtual chassis port: vcp-255/0/14
```

```

    Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                 : 4.130 mA
    Laser output power                 : 0.2450 mW / -6.11 dBm
    Module temperature                 : 32 degrees C / 90 degrees F
    Module voltage                     : 3.3530 V
    Receiver signal average optical power : 0.0971 mW / -10.13 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 : 14.998 mA
    Laser bias current low alarm threshold : 0.998 mA
    Laser bias current high warning threshold : 14.000 mA
    Laser bias current low warning threshold : 1.198 mA
    Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
    Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
    Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold : 85 degrees C / 185 degrees F
    Module temperature low alarm threshold : -10 degrees C / 14 degrees F
    Module temperature high warning threshold : 80 degrees C / 176 degrees F
    Module temperature low warning threshold : -5 degrees C / 23 degrees F
    Module voltage high alarm threshold : 3.600 V
    Module voltage low alarm threshold : 3.000 V
    Module voltage high warning threshold : 3.499 V

```



```

Module voltage low warning threshold      : 3.099 V
Laser rx power high alarm threshold       : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold        : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold     : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold      : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                       : 5.428 mA
Laser output power                       : 0.4760 mW / -3.22 dBm
Module temperature                       : 28 degrees C / 83 degrees F
Module voltage                           : 3.3440 V
Receiver signal average optical power    : 0.4002 mW / -3.98 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  : 10.500 mA
Laser bias current low alarm threshold   : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold  : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold    : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V

```

```

Module voltage low alarm threshold      : 2.970 V
Module voltage high warning threshold   : 3.465 V
Module voltage low warning threshold    : 3.135 V
Laser rx power high alarm threshold     : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

fpc3:

Virtual chassis port: vcp-255/0/2

```

Laser bias current                      : 7.876 mA
Laser output power                      : 0.5330 mW / -2.73 dBm
Module temperature                      : 26 degrees C / 78 degrees F
Module voltage                          : 3.3060 V
Receiver signal average optical power   : 0.4885 mW / -3.11 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 : 14.500 mA
Laser bias current low alarm threshold  : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm

```

```

Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold      : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                      : 5.052 mA
Laser output power                      : 0.5030 mW / -2.98 dBm
Module temperature                      : 24 degrees C / 75 degrees F
Module voltage                          : 3.2890 V
Receiver signal average optical power   : 0.5028 mW / -2.99 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                : Off
Laser rx power high warning             : Off
Laser rx power low warning              : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm

```

```

Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold    : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold        : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold      : 3.135 V
Laser rx power high alarm threshold       : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold        : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold      : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/4
Laser bias current                       : 7.978 mA
Laser output power                       : 0.5460 mW / -2.63 dBm
Module temperature                       : 24 degrees C / 76 degrees F
Module voltage                           : 3.3060 V
Receiver signal average optical power    : 0.6305 mW / -2.00 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  : 14.500 mA
Laser bias current low alarm threshold   : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA

```

```

Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold  : 75 degrees C / 167 degrees F
Module temperature low alarm threshold   : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold     : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

fpc4:

```

-----
Virtual chassis port: vcp-0
    Optical diagnostics : N/A
Virtual chassis port: vcp-1
    Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/4
    Laser bias current : 7.860 mA
    Laser output power : 0.5370 mW / -2.70 dBm
    Module temperature : 24 degrees C / 75 degrees F
    Module voltage     : 3.2920 V
    Receiver signal average optical power : 0.6271 mW / -2.03 dBm
    Laser bias current high alarm : Off
    Laser bias current low alarm  : Off
    Laser bias current high warning : Off
    Laser bias current low warning : Off
    Laser output power high alarm : Off
    Laser output power low alarm  : Off
    Laser output power high warning : Off
    Laser output power low warning : Off
    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 : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

show virtual-chassis vc-port diagnostics optics (interface-name)

user@external-routing-engine> **show virtual-chassis vc-port diagnostics optics vcp-255/0/3**

```

fpc0:
-----

fpc1:
-----

fpc2:
-----
Virtual chassis port: vcp-255/0/3
  Laser bias current           : 5.448 mA
  Laser output power           : 0.4770 mW / -3.21 dBm
  Module temperature           : 28 degrees C / 82 degrees F
  Module voltage               : 3.3450 V
  Receiver signal average optical power : 0.3973 mW / -4.01 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 : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold     : 3.630 V
Module voltage low alarm threshold      : 2.970 V
Module voltage high warning threshold   : 3.465 V
Module voltage low warning threshold    : 3.135 V
Laser rx power high alarm threshold     : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

fpc3:

Virtual chassis port: vcp-255/0/3

```

Laser bias current           : 5.040 mA

```

```

Laser output power           : 0.5020 mW / -2.99 dBm
Module temperature           : 24 degrees C / 74 degrees F
Module voltage                : 3.2870 V
Receiver signal average optical power : 0.5073 mW / -2.95 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 : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```



```
fpc4:
```

show virtual-chassis vc-port diagnostics optics local

```
user@switch> show virtual-chassis vc-port diagnostics optics local
```

```
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                 : 4.130 mA
  Laser output power                 : 0.2450 mW / -6.11 dBm
  Module temperature                 : 32 degrees C / 90 degrees F
  Module voltage                    : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 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 : 14.998 mA
  Laser bias current low alarm threshold : 0.998 mA
```

```

Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold  : 1.198 mA
Laser output power high alarm threshold   : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold     : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold   : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold    : 85 degrees C / 185 degrees F
Module temperature low alarm threshold     : -10 degrees C / 14 degrees F
Module temperature high warning threshold  : 80 degrees C / 176 degrees F
Module temperature low warning threshold   : -5 degrees C / 23 degrees F
Module voltage high alarm threshold        : 3.600 V
Module voltage low alarm threshold         : 3.000 V
Module voltage high warning threshold      : 3.499 V
Module voltage low warning threshold       : 3.099 V
Laser rx power high alarm threshold        : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold         : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold      : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold       : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                        : 5.426 mA
Laser output power                        : 0.4760 mW / -3.22 dBm
Module temperature                        : 28 degrees C / 83 degrees F
Module voltage                            : 3.3450 V
Receiver signal average optical power     : 0.3955 mW / -4.03 dBm
Laser bias current high alarm             : Off
Laser bias current low alarm              : Off
Laser bias current high warning           : Off
Laser bias current low warning            : Off
Laser output power high alarm             : Off
Laser output power low alarm              : Off
Laser output power high warning           : Off
Laser output power low warning            : Off
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 : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold      : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

show virtual-chassis vc-port diagnostics optics (member member-id)

user@switch> show virtual-chassis vc-port diagnostics optics member 2

```

fpc2:
-----
Virtual chassis port: vcp-2/0
    Optical diagnostics : N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current : 4.130 mA
    Laser output power : 0.2450 mW / -6.11 dBm
    Module temperature : 31 degrees C / 88 degrees F
    Module voltage     : 3.3530 V
    Receiver signal average optical power : 0.0961 mW / -10.17 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 : 14.998 mA
Laser bias current low alarm threshold : 0.998 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 1.198 mA
Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold : 85 degrees C / 185 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 80 degrees C / 176 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.600 V
Module voltage low alarm threshold : 3.000 V
Module voltage high warning threshold : 3.499 V
Module voltage low warning threshold : 3.099 V
Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current      : 5.418 mA
Laser output power      : 0.4770 mW / -3.21 dBm
Module temperature      : 28 degrees C / 83 degrees F
Module voltage          : 3.3450 V
Receiver signal average optical power : 0.3964 mW / -4.02 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 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 : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold     : 3.630 V
Module voltage low alarm threshold      : 2.970 V
Module voltage high warning threshold   : 3.465 V
Module voltage low warning threshold    : 3.135 V
Laser rx power high alarm threshold     : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

test interface restart-auto-negotiation

Syntax

```
test interface restart-auto-negotiation interface-name
```

Release Information

Command introduced in Junos OS Release 7.6.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Restarts auto-negotiation on a Fast Ethernet or Gigabit Ethernet interface.

Options

interface-name—Interface name: ***fe-fpc/pic/port*** or ***ge-fpc/pic/port***.

Required Privilege Level

view

List of Sample Output

[test interface restart-auto-negotiation on page 1334](#)

Output Fields

Use the **show interfaces extensive** command to see the state for auto-negotiation.

Sample Output

test interface restart-auto-negotiation

```
user@host> test interface restart-auto-negotiation fe-1/0/0
```