JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Interfaces User Guide for Security Devices

Published
2022-06-13

JUNOS

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

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

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

Junos® OS Interfaces User Guide for Security Devices
Copyright © 2022 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About This Guide | xxiii

1 **Overview**

Introduction to Interfaces | 2

- Understanding Interfaces | 2
- Network Interfaces | 3
- Services Interfaces | 5
- Special Interfaces | 8
- Interface Naming Conventions | 9
- Understanding the Data Link Layer | 12

Physical Interface Properties | 14

- Understanding Interface Physical Properties | 15
- Understanding Bit Error Rate Testing | 17
- Understanding Interface Clocking | 18
- Understanding Frame Check Sequences | 19
- MTU Default and Maximum Values | 20
- Understanding Jumbo Frames Support for Ethernet Interfaces | 24

Logical Interface Properties | 25

- Understanding Interface Logical Properties | 25
- Understanding Protocol Families | 26

Understanding IPv4 and IPv6 Protocol Family | 27

- Understanding IPv4 Addressing | 28
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types | 32
- Configuring the inet6 IPv6 Protocol Family | 36

Configuring VLAN Tagging | 38

Understanding Virtual LANs | 38

VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices | 40

Configuring VLAN Tagging | 41

2 Configuring DS1, DS3, and 1-Port Clear Channel DS3/E3 GPIM Interfaces

Configuring DS1 Interfaces | 47

Understanding T1 and E1 Interfaces | 47

Example: Configuring a T1 Interface | 51

Requirements | 51

Overview | 51

Configuration | 52

Verification | 53

Example: Deleting a T1 Interface | 55

Requirements | 55

Overview | 55

Configuration | 56

Verification | 56

Configuring DS3 Interfaces | 57

Understanding T3 and E3 Interfaces | 57

Example: Configuring a T3 Interface | 63

Requirements | 63

Overview | 63

Configuration | 64

Verification | 65

Example: Deleting a T3 Interface | 67

Requirements | 67

Overview | 67

Configuration | 68

Verification | 68

Configuring 1-Port Clear Channel DS3/E3 GPIM | 68

Understanding the 1-Port Clear Channel DS3/E3 GPIM | 69

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode | 73

Requirements | 73
Overview | 74
Configuration | 74

Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode | 76

Requirements | 76
Overview | 76
Configuration | 76

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode | 78

Requirements | 78
Overview | 78
Configuration | 79

3 Configuring ADSL and SHDSL Interfaces

ADSL and SHDSL Interfaces | 82

ADSL and SHDSL Interface Overview | 82

Example: Configure ADSL and SHDSL Network Interfaces | 86

Example: Configure G.SHDSL Interface | 106

VDSL2 Interfaces | 125

VDSL2 Interface Overview | 125

Example: Configure VDSL2 Interface | 129

Configure the VDSL2 Interface and Enable VLAN Tagging | 133
Configure VDSL2 Interface with VDSL2 Mini-PIMs | 135
Verification | 142

4 Configuring Ethernet Interfaces

Ethernet Interfaces | 160

Ethernet Interfaces Overview | 160

Example: Configure Ethernet Interface | 166

Overview | 166

Example: Configuring Promiscuous Mode on the SRX5K-MPC | 167

Verification | 168

Configuring Aggregated Ethernet Interfaces | 171

Understanding Aggregated Ethernet Interfaces | **172**

Configuring Aggregated Ethernet Interfaces | **174**

Understanding Physical Interfaces for Aggregated Ethernet Interfaces | **175**

Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces | **176**

- Requirements | **176**
- Overview | **176**
- Configuration | **176**
- Verification | **177**

Understanding Aggregated Ethernet Interface Link Speed | **177**

Example: Configuring Aggregated Ethernet Link Speed | **178**

- Requirements | **178**
- Overview | **178**
- Configuration | **179**
- Verification | **179**

Understanding Minimum Links for Aggregated Ethernet Interfaces | **179**

Example: Configuring Aggregated Ethernet Minimum Links | **180**

- Requirements | **180**
- Overview | **180**
- Configuration | **181**
- Verification | **181**

Deleting Aggregated Ethernet Interface | **181**

Example: Deleting Aggregated Ethernet Interfaces | **182**

- Requirements | **182**
- Overview | **182**
- Configuration | **182**
- Verification | **183**

Example: Deleting Aggregated Ethernet Interface Contents | **183**

- Requirements | **183**
- Overview | **184**
- Configuration | **184**
- Verification | **184**

Understanding VLAN Tagging for Aggregated Ethernet Interfaces | **185**

Understanding Promiscuous Mode for Aggregated Ethernet Interfaces | **185**

Verifying Aggregated Ethernet Interfaces | **185**

- Verifying Aggregated Ethernet Interfaces (terse) | **186**
- Verifying Aggregated Ethernet Interfaces (extensive) | **187**

Configuring Link Aggregation Control Protocol | 188

Understanding LACP on Standalone Devices | **189**

Example: Configuring Link Aggregation Control Protocol | **189**

- Requirements | **190**
- Overview | **190**
- Configuration | **190**
- Verification | **193**

Verifying LACP on Standalone Devices | **195**

- Verifying LACP Statistics | **195**
- Verifying LACP Aggregated Ethernet Interfaces | **196**

LAG and LACP Support Line Devices with I/O Cards (IOCs) | **198**

Example: Configuring LAG Interface on an Line Device with IOC2 or IOC3 | **200**

- Requirements | **200**
- Overview | **201**
- Configuration | **201**
- Verification | **205**

Configuring Gigabit Ethernet Physical Interface Modules | 207

Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM | **207**

Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface | **210**

- Requirements | **210**
- Overview | **210**
- Configuration | **211**
- Verification | **215**

Understanding the 2-Port 10-Gigabit Ethernet XPIM | **219**

Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface | **222**

Requirements | 222
Overview | 222
Configuration | 223
Verification | 225

Understanding the 8-Port Gigabit Ethernet SFP XPIM | 228

Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs | 230
Requirements | 231
Overview and Topology | 231
Configuration | 232
Verification | 238

Port Speed on SRX Series Devices | 251

SRX4600 Port Speed Overview | 251

Port Speed on SRX5K-IOC4-MRATE | 256
Configuring Port Speed at PIC Level | 256
Configuring Port Speed at Port Level | 258

Power over Ethernet | 261

Power over Ethernet Overview | 261

Example: Configure PoE Interface | 268
Verification | 271

5 **Configuring Interface Encapsulation**

Interface Encapsulation Overview | 274

Understanding Physical Encapsulation on an Interface | 274

Understanding Frame Relay Encapsulation on an Interface | 275

Understanding Point-to-Point Protocol | 277

Understanding High-Level Data Link Control | 280

Configuring GRE Keepalive Time | 281

Understanding GRE Keepalive Time | 282

Configuring GRE Keepalive Time | 283
Configuring Keepalive Time and Hold time for a GRE Tunnel Interface | 283
Display GRE Keepalive Time Configuration | 284

Display Keepalive Time Information on a GRE Tunnel Interface | **285**

Example: GRE Configuration | **288**

Requirements | **288**

Overview | **288**

Configuration | **288**

Verification | **292**

Example: Configuring GRE over IPsec Tunnels | **295**

Requirements | **295**

Overview | **295**

Configuration | **296**

Verification | **299**

Example: Configuring a GRE Tunnel When the Tunnel Destination Is in a Routing Instance | **300**

Requirements | **300**

Overview | **300**

Configuration | **301**

Verification | **306**

Configuring Point-to-Point Protocol over Ethernet | 307

Understanding Point-to-Point Protocol over Ethernet | **308**

Understanding PPPoE Interfaces | **312**

Example: Configuring PPPoE Interfaces | **312**

Requirements | **312**

Overview | **313**

Configuration | **313**

Disabling the End-of-List Tag | **319**

Understanding PPPoE Ethernet Interfaces | **322**

Example: Configuring PPPoE Encapsulation on an Ethernet Interface | **322**

Requirements | **322**

Overview | **323**

Configuration | **323**

Verification | **323**

Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces | **324**

Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface | **324**

- Requirements | **325**
- Overview | **325**
- Configuration | **325**
- Verification | **327**

Understanding CHAP Authentication on a PPPoE Interface | **328**

Example: Configuring CHAP Authentication on a PPPoE Interface | **328**

- Requirements | **328**
- Overview | **329**
- Configuration | **329**
- Verification | **331**

Verifying Credit-Flow Control | **331**

Verifying PPPoE Interfaces | **333**

Verifying R2CP Interfaces | **334**

Displaying Statistics for PPPoE | **336**

Setting Tracing Options for PPPoE | **337**

6 **Configuring Link Services Interfaces**

Configuring Link Services Interfaces | 340

Link Services Interfaces Overview | **340**

Link Services Configuration Overview | **348**

Verifying the Link Services Interface | **349**

- Verifying Link Services Interface Statistics | **349**
- Verifying Link Services CoS Configuration | **352**

Understanding the Internal Interface LSQ-0/0/0 Configuration | **354**

Example: Upgrading from ls-0/0/0 to lsq-0/0/0 for Multilink Services | **355**

- Requirements | **355**
- Overview | **355**
- Configuration | **356**
- Verification | **359**

Troubleshooting the Link Services Interface | **359**

Determine Which CoS Components Are Applied to the Constituent Links | **360**

Determine What Causes Jitter and Latency on the Multilink Bundle | **362**

Determine If LFI and Load Balancing Are Working Correctly | **363**

Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device | **372**

Configuring Link Fragmentation and Interleaving | 372

Understanding Link Fragmentation and Interleaving Configuration | **373**

Example: Configuring Link Fragmentation and Interleaving | **374**

Requirements | **374**

Overview | **375**

Configuration | **375**

Verification | **376**

Configuring Class-of-Service on Link Services Interfaces | 376

Understanding How to Define Classifiers and Forwarding Classes | **377**

Example: Defining Classifiers and Forwarding Classes | **377**

Requirements | **378**

Overview | **378**

Configuration | **378**

Verification | **381**

Understanding How to Define and Apply Scheduler Maps | **382**

Example: Configuring Scheduler Maps | **384**

Requirements | **384**

Overview | **384**

Configuration | **385**

Verification | **388**

Understanding Interface Shaping Rates | **389**

Example: Configuring Interface Shaping Rates | **389**

Requirements | **389**

Overview | **390**

Configuration | **390**

Verification | **391**

Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles | 391

Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | 391

Example: Configuring an MLPPP Bundle | 392

Requirements | 393

Overview | 393

Configuration | 393

Verification | 397

Configuring Multilink Frame Relay | 397

Understanding Multilink Frame Relay FRF.15 | 398

Example: Configuring Multilink Frame Relay FRF.15 | 398

Requirements | 398

Overview | 398

Configuration | 399

Verification | 402

Understanding Multilink Frame Relay FRF.16 | 402

Example: Configuring Multilink Frame Relay FRF.16 | 403

Requirements | 403

Overview | 403

Configuration | 404

Verification | 409

Configuring Compressed Real-Time Transport Protocol | 409

Understanding Compressed Real-Time Transport Protocol | 410

Example: Configuring the Compressed Real-Time Transport Protocol | 410

Requirements | 410

Overview | 411

Configuration | 411

Verification | 413

7 **Configuring Management, Discard, and Loopback Interfaces**

Configuring Management and Discard Interfaces | 415

Configuring Management Interfaces | 415

Configuring Discard Interface | 416

Configuring Loopback Interfaces | 416

Loopback Interface Overview | 416

Configuring a Loopback Interface | 417

8 LTE Mini-PIM

LTE Mini Physical Interface Modules (LTE Mini-PIM) | 421

LTE Mini-PIM Overview | 421

Configure LTE Mini-PIM | 424

Configure LTE Mini-PIM as a Primary Interface | 424

Configure LTE Mini-PIM in a High Availability Cluster Mode | 426

Configure LTE Mini-PIM as a Backup Interface | 428

Configure LTE Mini-PIM as a Dial-on-demand Interface | 430

Example: Configure LTE Mini-PIM as a Backup Interface | 433

Requirements | 433

Overview | 433

Configuration | 433

Verification | 436

9 Wi-Fi MPIM

Wi-Fi Mini Physical Interface Module (MPIM) | 443

Wi-Fi Mini-Physical Interface Module Overview | 443

Configure Wi-Fi Mini-PIM | 446

Configure Network Setting for the Wi-Fi Mini-PIM | 447

Configure VLANS | 452

Configure Multiple VLANs and SSIDs | 454

10 Interfaces Support for SRX100, SRX110, SRX210, SRX240, SRX550, SRX650, and SRX1400 Devices

Configuring 1-Port Clear Channel DS3/E3 GPIM | 461

Understanding the 1-Port Clear Channel DS3/E3 GPIM | 461

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode | 466

Requirements | 466

Overview | 466

Configuration | 467

Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode | 468

Requirements | 469
Overview | 469
Configuration | 469

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode | 471
Requirements | 471
Overview | 471
Configuration | 472

Configuring 3G Wireless Modems for WAN Connections | 473

3G Wireless Modem Overview | 474

3G Wireless Modem Configuration Overview | 475

Understanding the Dialer Interface | 477

Example: Configuring the Dialer Interface | 479
Requirements | 479
Overview | 480
Configuration | 480
Verification | 487

Understanding the 3G Wireless Modem Physical Interface | 488

Example: Configuring the 3G Wireless Modem Interface | 488
Requirements | 489
Overview | 489
Configuration | 489
Verification | 490

Understanding the GSM Profile | 490

Example: Configuring the GSM Profile | 491
Requirements | 491
Overview | 491
Configuration | 492
Verification | 493

Unlocking the GSM 3G Wireless Modem | 493

Configuring CDMA EV-DO Modem Cards | 494

Understanding Account Activation for CDMA EV-DO Modem Cards | 495

Activating the CDMA EV-DO Modem Card Manually | **497**

Activating the CDMA EV-DO Modem Card with IOTA Provisioning | **499**

Activating the CDMA EV-DO Modem Card with OTASP Provisioning | **499**

Configuring USB Modems for Dial Backup | 501

USB Modem Interface Overview | **501**

USB Modem Configuration Overview | **505**

Example: Configuring a USB Modem Interface | **507**

Requirements | **508**

Overview | **508**

Configuration | **508**

Verification | **510**

Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup | **511**

Requirements | **512**

Overview | **512**

Configuration | **512**

Verification | **521**

Example: Configuring a Dialer Interface for USB Modem Dial-In | **521**

Requirements | **522**

Overview | **522**

Configuration | **523**

Verification | **524**

Example: Configuring PAP on Dialer Interfaces | **524**

Requirements | **524**

Overview | **524**

Configuration | **524**

Verification | **525**

Example: Configuring CHAP on Dialer Interfaces | **525**

Requirements | **526**

Overview | **526**

Configuration | **526**

Verification | **527**

Configuring DOCSIS Mini-PIM Interfaces | 527

DOCSIS Mini-PIM Interface Overview | **528**

Software Features Supported on DOCSIS Mini-PIMs | **530**

Example: Configuring the DOCSIS Mini-PIM Interfaces | **531**

Requirements | **532**

Overview | **532**

Configuration | **532**

Verification | **534**

11 **Configuration Statements**

accept-source-mac | 543

access-point name | 545

apply-groups | 546

activation-delay | 548

authentication-method (Interfaces) | 549

bandwidth (Interfaces) | 551

bundle (Interfaces) | 552

cbr rate | 553

callback | 555

callback-wait-period | 556

caller | 558

cellular-options | 560

classifiers (CoS) | 561

client-identifier (Interfaces) | 563

code-points (CoS) | 565

compression-device (Interfaces) | 567

credit (Interfaces) | 568

data-rate | 569

deactivation-delay | 571

disable (PoE) | 573

dialer-options | 574

dialin | 576

dial-string | 578

dhcp (DHCP Client) | 579

dsl-sfp-options | 582

duration (PoE) | 586

family inet (Interfaces) | 587

family inet6 | 592

flag (Interfaces) | 596

flexible-vlan-tagging (Interfaces) | 598

flow-control (Interfaces) | 599

flow-monitoring (Services) | 601

forwarding-classes | 603

fpc (Interfaces) | 609

gratuitous-arp-reply | 610

gsm-options | 612

guard-band (PoE) | 614

hold-time (Redundant Ethernet Interfaces) | 615

hub-assist | 617

idle-timeout | 619

incoming-map | 620

initial-route-check | 622

inline-jflow (Forwarding Options) | 623

interface (PIC Bundle) | 625

interface (PoE) | 627

interfaces (CoS) | 629

interval (Interfaces) | 631

interval (PoE) | 633

isdn-options | 634

ipv4-template (Services) | 636

ipv6-template (Services) | 637

lacp (Interfaces) | 639

latency (Interfaces) | 641

lease-time | 642

line-rate (Interfaces) | 644

link-speed (Interfaces) | 645

load-interval | 647

load-threshold | 648

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | 650

loss-priority (CoS Loss Priority) | 652

loss-priority (CoS Rewrite Rules) | 654

loss-priority-maps (CoS Interfaces) | 656

loss-priority-maps (CoS) | 657

management (PoE) | 659

maximum-power (PoE) | 661

mdi-mode | 662

media-type (Interfaces) | 665

minimum-links (Interfaces) | 667

modem-options | 668

mtu | 670

native-vlan-id | 671

next-hop-tunnel | 675

no-dns-propagation | 677

option-refresh-rate (Services) | 678

pic-mode (Chassis T1 Mode) | 680

periodic (Interfaces) | 681

pool | 683

ppp-over-ether | 685

pppoe | 686

pppoe-options | 688

priority (PoE) | 690

profile (Access) | 692

profiles | 697

promiscuous-mode (Interfaces) | 698

quality (Interfaces) | 700

r2cp | 701

radio-router (Interfaces) | 703

redial-delay | 705

redundancy-group (Interfaces) | 707

redundant-ether-options | 708

redundant-parent (Interfaces Fast Ethernet) | 711

redundant-parent (Interfaces Gigabit Ethernet) | 712

request pppoe connect | 714

request pppoe disconnect | 715

resource (Interfaces) | 717

retransmission-attempt (DHCP Client) | 719

retransmission-interval (DHCP Client) | 721

roaming-mode | 722

scheduler-map (CoS Virtual Channels) | 724

select-profile | 725

server-address | 727

shaping-rate (CoS Interfaces) | 729

simple-filter (Interfaces) | 731

sip-password | 732

sip-user-id | 734

source-address-filter (Interfaces) | 735

source-filtering (Interfaces) | 737

speed (Interfaces) | 739

speed (Gigabit Ethernet interface) | 740

spid1 | 742

spid2 | 743

static-tei-val | 745

switch-type | 746

t310 | 748

tei-option | 749

telemetries (PoE) | 751

template-refresh-rate (Services) | 753

threshold (Interfaces) | 754

traceoptions (Interfaces) | 756

update-server | 757

vbr rate | 759

vdsl-profile | 760

vendor-id (Interfaces) | 762

watch-list | 764

web-authentication (Interfaces) | 765

wlan | 767

12 Operational Commands

clear oam ethernet connectivity-fault-management path-database | 775

clear dhcpv6 server binding (Local Server) | 776

clear ethernet-switching statistics mac-learning | 778

clear interfaces statistics swfabx | 780

clear ipv6 neighbors | 781

clear lacp statistics interfaces | 783

restart | 785

request modem wireless create-profile | 802

request modem wireless fota | 804

request modem wireless sim-lock | 806

request modem wireless sim-unlock | 808

request wlan access-point packet capture | 810

show chassis fpc (View) | 812

show chassis hardware (View) | 824

show ethernet-switching mac-learning-log | 845

show ethernet-switching table | 851

show igmp-snooping route (View) | 885

show interfaces | 887

show interfaces diagnostics optics | 1023

show interfaces flow-statistics | 1030

show interfaces queue | 1037

show interfaces statistics (View) | 1044

show interfaces terse zone | 1046

show ipv6 neighbors | 1047

show lacp interfaces (View) | 1050

show lacp statistics interfaces (View) | 1056

show modem wireless firmware | 1059

show modem wireless network | 1063

show modem wireless profiles | 1068

show oam ethernet link-fault-management | 1071

show poe controller (View) | 1080

show pppoe interfaces | 1082

show pppoe statistics | 1088

show poe telemetries | 1092

show services accounting | 1095

show services accounting aggregation (View) | 1099

show services accounting aggregation template (View) | 1100

show services accounting flow-detail (View) | 1102

show wlan access-points | 1103

speed (Chassis Cluster) | 1112

About This Guide

Use this guide to configure and monitor Network, Services, and Special interfaces for Juniper security devices.

- Refer to LTE interfaces and Wi-Fi Mini-PIM interfaces on SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550 HM devices.

  Also, understand and configure the physical, logical and VLAN interfaces, DS1 and DS3 interfaces, ADSL, SHDSL, and VDSL interfaces, Ethernet Interfaces, interface encapsulation, link service interfaces, management, discard, and loopback interfaces, and serial interfaces on SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550 HM devices.

- Refer to Interfaces Support for SRX100, SRX110, SRX210, SRX240, SRX550, SRX650, and SRX1400 Devices section to access information on modem interfaces and 1-Port Clear Channel DS3/E3 GPIM interfaces.

- Refer to Interfaces Fundamentals for information on serial interfaces.

1

CHAPTER

Overview

Introduction to Interfaces | 2

Physical Interface Properties | 14

Logical Interface Properties | 25

Understanding IPv4 and IPv6 Protocol Family | 27

Configuring VLAN Tagging | 38

Introduction to Interfaces

IN THIS SECTION

- Understanding Interfaces | 2
- Network Interfaces | 3
- Services Interfaces | 5
- Special Interfaces | 8
- Interface Naming Conventions | 9
- Understanding the Data Link Layer | 12

Junos OS supports different types of interfaces on which the devices function. The following topics provide information of types of interfaces used on security devices, the naming conventions and how to monitor the interfaces.

Understanding Interfaces

Interfaces act as a doorway through which traffic enters and exits a device. Juniper Networks devices support a variety of interface types:

- Network interfaces—Networking interfaces primarily provide traffic connectivity.
- Services interfaces—Services interfaces manipulate traffic before it is delivered to its destination.
- Special interfaces—Special interfaces include management interfaces, the loopback interface, and the discard interface.

Each type of interface uses a particular medium to transmit data. The physical wires and Data Link Layer protocols used by a medium determine how traffic is sent. To configure and monitor interfaces, you need to understand their media characteristics, as well as physical and logical properties such as IP addressing, link-layer protocols, and link encapsulation.

NOTE: Most interfaces are configurable, but some internally generated interfaces are not configurable.

Network Interfaces

All Juniper Networks devices use network interfaces to make physical connections to other devices. A connection takes place along media-specific physical wires through an I/O card (IOC) in the SRX Series Services Gateway. Networking interfaces primarily provide traffic connectivity.

You must configure each network interface before it can operate on the device. Configuring an interface can define both the physical properties of the link and the logical properties of a logical interface on the link.

Table 1 on page 3 describes network interfaces that are available on SRX Series devices.

Table 1: Network Interfaces

Interface Name	Description
ae	Aggregated Ethernet interface. See "Understanding Aggregated Ethernet Interfaces" on page 172.
at	ATM-over-ADSL or ATM-over-SHDSL WAN interface.
cl	Physical interface for the 3G wireless modem or LTE Mini-PIM. See "Understanding the 3G Wireless Modem Physical Interface" on page 488 and LTE Mini-PIM Overview. Starting with Junos OS Release 15.1X49-D100, SRX320, SRX340, SRX345, and SRX550HM devices support the LTE interface. The dialer interface is used for initiating wireless WAN connections over LTE networks.
dl	Dialer interface for initiating USB modem or wireless WAN connections. See *USB Modem Interface Overview* and LTE Mini-PIM Overview.
e1	E1 (also called DS1) WAN interface. See "Understanding T1 and E1 Interfaces" on page 47.

Table 1: Network Interfaces *(Continued)*

Interface Name	Description
e3	E3 (also called DS3) WAN interface. See "Understanding T3 and E3 Interfaces" on page 57.
fe	Fast Ethernet interface. See "Understanding Ethernet Interfaces" on page 160.
ge	Gigabit Ethernet interface. See "Understanding Ethernet Interfaces" on page 160.
pt	VDSL2 interface. See Example: Configuring VDSL2 Interfaces (Detail).
reth	For chassis cluster configurations only, redundant Ethernet interface. See "Understanding Ethernet Interfaces" on page 160.
se	Serial interface (either RS-232, RS-422/499, RS-530, V.35, or X.21). See Serial Interfaces Overview.
t1	T1 (also called DS1) WAN interface. See "Understanding T1 and E1 Interfaces" on page 47.
t3	T3 (also called DS3) WAN interface. See "Understanding T3 and E3 Interfaces" on page 57.
wx	WXC Integrated Services Module (ISM 200) interface for WAN acceleration. See the WXC Integrated Services Module Installation and Configuration.
xe	10-Gigabit Ethernet interface. See "Understanding the 2-Port 10-Gigabit Ethernet XPIM" on page 219.

NOTE: The affected interfaces are these: ATM-over-ADSL or ATM-over-SHDSL (at) interface, dialer interface (dl), E1 (also called DS1) WAN interface, E3 (also called DS3) WAN interface, VDSL2 interface (pt), serial interface (se), T1 (also called DS1) WAN interface, T3 (also called DS3) WAN interface. However, starting from Junos OS Release 15.1X49-D40 and onwards, SRX300,

SRX320, SRX340, SRX345, SRX380, and SRX550HM devices support VDSL2 (pt), serial (se), T1 (t1) , and E1 (e1) interfaces.

Services Interfaces

Services interfaces provide specific capabilities for manipulating traffic before it is delivered to its destination. On Juniper Networks M Series and T Series routing platforms, individual services such as IP-over-IP encapsulation, link services such as multilink protocols, adaptive services such as stateful firewall filters and NAT, and sampling and logging capabilities are implemented by services Physical Interface Cards (PICs). On SRX Series devices, services processing is handled by the Services Processing Card (SPC).

Although the same Junos OS image supports the services features across all routing platforms, on SRX Series devices, services interfaces are not associated with a physical interface. To configure services on these devices, you configure one or more internal interfaces by specifying slot `0`, interface carrier `0`, and port `0`—for example, `gr-0/0/0` for GRE.

Table 2 on page 5 describes services interfaces that you can configure on SRX Series Services Gateways.

Table 2: Configurable Services Interfaces

Interface Name	Description
`gr-0/0/0`	Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol inside another routing protocol. Packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then sent. You can create multiple instances of this interface for forwarding encapsulated data to multiple destination addresses by using the default interface as the parent and creating extensions, for example, gr-0/0/0.1, gr-0/0/0.2, and so on. The GRE interface is an internal interface only and is not associated with a physical interface. It is used only for processing GRE traffic. See the Junos OS Services Interfaces Library for Routing Devices for information about tunnel services.

Table 2: Configurable Services Interfaces *(Continued)*

Interface Name	Description
ip-0/0/0	Configurable IP-over-IP encapsulation (IP-IP tunnel) interface. IP tunneling allows the encapsulation of one IP packet inside another IP packet. With IP routing, you can route IP packets directly to a particular address or route the IP packets to an internal interface where they are encapsulated inside an IP-IP tunnel and forwarded to the encapsulating packet's destination address. You can create multiple instances of this interface for forwarding IP-IP tunnel data to multiple destination addresses by using the default interface as the parent and creating extensions, for example, ip-0/0/0.1, ip-0/0/0.2, and so on. The IP-IP interface is an internal interface only and is not associated with a physical interface. It is used only for processing IP-IP tunnel traffic. See the Junos OS Services Interfaces Library for Routing Devices for information about tunnel services.
lsq-0/0/0	Configurable link services queuing interface. Link services include the multilink services MLPPP, MLFR, and Compressed Real-Time Transport Protocol (CRTP). Packets are routed to this internal interface for link bundling or compression. The link services interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform multilink services. **NOTE**: The ls-0/0/0 interface has been deprecated. All multiclass multilink features supported by ls-0/0/0 are now supported by lsq-0/0/0.
lt-0/0/0	Configurable logical tunnel interface that interconnects logical systems on SRX Series devices. See the Logical Systems and Tenant Systems User Guide for Security Devices.
pp0	Configurable PPPoE encapsulation interface. PPP packets being routed in an Ethernet network use PPPoE encapsulation. Packets are routed to this internal interface for PPPoE encapsulation. The PPPoE encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to forward PPPoE traffic. See "Understanding Point-to-Point Protocol over Ethernet" on page 308.

Table 2: Configurable Services Interfaces *(Continued)*

Interface Name	Description
ppd0	Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point device. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree. Within a device, packets are routed to this internal interface for de-encapsulation. The PIM de-encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure PIM with the `[edit protocol pim]` hierarchy to perform PIM de-encapsulation. Use the `show pim interfaces` command to check the status of ppd0 interface.
ppe0	Protocol Independent Multicast (PIM) encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point device. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree. Within a device, packets are routed to this internal interface for encapsulation. The PIM encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure PIM with the `[edit protocol pim]` hierarchy to perform PIM encapsulation.
st0	Secure tunnel interface used for IPSec VPNs. See the IPsec VPN User Guide for Security Devices.
umd0	Configurable USB modem physical interface. This interface is detected when a USB modem is connected to the USB port on the device. See *USB Modem Configuration Overview*.
mt-0/0/0	Multicast tunnel interface. This interface is automatically generated, but you can configure properties on it if needed.

Table 3 on page 8 describes non-configurable services interfaces for SRX Series Services Gateways.

Table 3: Non-Configurable Services Interfaces

Interface Name	Description
`gre`	Internally generated Generic Routing Encapsulation (GRE) interface created by Junos OS to handle GRE traffic. It is not a configurable interface.
`ipip`	Internally generated IP-over-IP interface created by Junos OS to handle IP tunnel traffic. It is not a configurable interface.
`lsi`	Internally generated link services interface created by Junos OS to handle multilink services like MLPPP, MLFR, and CRTP. It is not a configurable interface.
`pc-`*`pim`*`/0/0`	Internally configured interface used by the system as a control path between the WXC Integrated Services Module and the Routing Engine. It is not a configurable interface. See the WX and WXC Series.
`pimd`	Internally generated Protocol Independent Multicast (PIM) de-encapsulation interface created by Junos OS to handle PIM de-encapsulation. It is not a configurable interface.
`pime`	Internally generated Protocol Independent Multicast (PIM) encapsulation interface created by Junos OS to handle PIM encapsulation. It is not a configurable interface.
`tap`	Internally generated interface created by Junos OS to monitor and record traffic during passive monitoring. Packets discarded by the Packet Forwarding Engine are placed on this interface. It is not a configurable interface.
`sp-0/0/0`	Adaptive services interface. The logical interface `sp-`*`fpc/pic/port`*. 16383 is an internally generated, non-configurable interface for router control traffic.

Special Interfaces

Special interfaces include management interfaces, which are primarily intended for accessing the device remotely, the loopback interface, which has several uses depending on the particular Junos OS feature being configured, and the discard interface.

Table 4 on page 9 describes special interfaces for SRX Series Services Gateways.

Table 4: Special Interfaces

Interface Name	Description
`fxp0, fxp1`	On SRX Series devices, the fxp0 management interface is a dedicated port located on the Routing Engine.
`lo0`	Loopback address. The loopback address has several uses, depending on the particular Junos feature being configured.
`dsc`	Discard interface.

Interface Naming Conventions

Each device interface has a unique name that follows a naming convention. If you are familiar with Juniper Networks M Series and T Series routing platforms, be aware that device interface names are similar to but not identical to the interface names on those routing platforms.

The unique name of each network interface identifies its type and location and indicates whether it is a physical interface or an optional logical unit created on a physical interface.

- The name of each network interface has the following format to identify the physical device that corresponds to a single physical network connector:

  ```
  type-slot/pim-or-ioc/port
  ```

- Network interfaces that are fractionalized into time slots include a channel number in the name, preceded by a colon (:):

  ```
  type-slot/pim-or-ioc/port:channel
  ```

- Each logical interface has an additional logical unit identifier, preceded by a period (.):

  ```
  type-slot/pim-or-ioc/port:<channel>.unit
  ```

The parts of an interface name are summarized in Table 5 on page 10.

Table 5: Network Interface Names

Name Part	Meaning	Possible Values
type	Type of network medium that can connect to this interface.	ae, at, ei, e3, fe, fxp0, fxp1, ge, lo0, lsq, lt, ppo, pt, sto, t1, t3, xe, and so on.
slot	Number of the chassis slot in which a PIM or IOC is installed.	SRX5600 and SRX5800 devices: The slot number begins at `0` and increases as follows from left to right, bottom to top: • SRX5600 device—Slots 0 to 5 • SRX5800 device—Slots 0 to 5, 7 to 11 SRX3400 and SRX3600 devices: The Switch Fabric Board (SFB) is always `0`. Slot numbers increase as follows from top to bottom, left to right: • SRX3400 devce—Slots 0 to 4 • SRX3600 device—Slots 0 to 6 • SRX4600 device—Slots 0 to 6
pim-or-ioc	Number of the PIM or IOC on which the physical interface is located.	SRX5600 and SRX5800 devices: For 40-port Gigabit Ethernet IOCs or 4-port 10-Gigabit Ethernet IOCs, this number can be `0`, `1`, `2`, or `3`. SRX3400, SRX3600, and SRX 4600 devices: This number is always `0`. Only one IOC can be installed in a slot.

Table 5: Network Interface Names *(Continued)*

Name Part	Meaning	Possible Values
port	Number of the port on a PIM or IOC on which the physical interface is located.	On SRX5600 and SRX5800 devices: • For 40-port Gigabit Ethernet IOCs, this number begins at 0 and increases from left to right to a maximum of 9. • For 4-port 10-Gigabit Ethernet IOCs, this number is always 0. On SRX3400, SRX3600, and SRX 4600 devices: • For the SFB built-in copper Gigabit Ethernet ports, this number begins at 0 and increases from top to bottom, left to right, to a maximum of 7. For the SFB built-in fiber Gigabit Ethernet ports, this number begins at 8 and increases from left to right to a maximum of 11. • For 16-port Gigabit Ethernet IOCs, this number begins at 0 to a maximum of 15. • For 2-port 10-Gigabit Ethernet IOCs, this number is 0 or 1. Port numbers appear on the PIM or IOC faceplate.
channel	Number of the channel (time slot) on a fractional or channelized T1 or E1 interface.	• On an E1 interface, a value from 1 through 31. The 1 time slot is reserved. • On a T1 interface, a value from 1 through 24.
unit	Number of the logical interface created on a physical interface.	A value from 0 through 16384. If no logical interface number is specified, unit 0 is the default, but must be explicitly configured. In addition to user-configured interfaces, there are some logical interfaces that are created dynamically. Hence, for Junos OS, the maximum limit for configuring logical interfaces is 2,62,143 (user configured and dynamically created). Based on performance, for each platform, the maximum number of logical interfaces supported can vary.

NOTE: Platform support depends on the Junos OS release in your installation.

Understanding the Data Link Layer

IN THIS SECTION

- Physical Addressing | 12
- Network Topology | 12
- Error Notification | 13
- Frame Sequencing | 13
- Flow Control | 13
- Data Link Sublayers | 13
- MAC Addressing | 13

The Data Link Layer is Layer 2 in the Open Systems Interconnection (OSI) model. The Data Link Layer is responsible for transmitting data across a physical network link. Each physical medium has link-layer specifications for network and link-layer protocol characteristics such as physical addressing, network topology, error notification, frame sequencing, and flow control.

Physical Addressing

Physical addressing is different from network addressing. Network addresses differentiate between nodes or devices in a network, allowing traffic to be routed or switched through the network. In contrast, physical addressing identifies devices at the link-layer level, differentiating between individual devices on the same physical medium. The primary form of physical addressing is the media access control (MAC) address.

Network Topology

Network topology specifications identify how devices are linked in a network. Some media allow devices to be connected by a bus topology, while others require a ring topology. The bus topology is used by Ethernet technologies, which are supported on Juniper Networks devices.

Error Notification

The Data Link Layer provides error notifications that alert higher layer protocols that an error has occurred on the physical link. Examples of link-level errors include the loss of a signal, the loss of a clocking signal across serial connections, or the loss of the remote endpoint on a T1 or T3 link.

Frame Sequencing

The frame sequencing capabilities of the Data Link Layer allow frames that are transmitted out of sequence to be reordered on the receiving end of a transmission. The integrity of the packet can then be verified by means of the bits in the Layer 2 header, which is transmitted along with the data payload.

Flow Control

Flow control within the Data Link Layer allows receiving devices on a link to detect congestion and notify their upstream and downstream neighbors. The neighbor devices relay the congestion information to their higher layer protocols so that the flow of traffic can be altered or rerouted.

Data Link Sublayers

The Data Link Layer is divided into two sublayers: logical link control (LLC) and media access control (MAC). The LLC sublayer manages communications between devices over a single link of a network. This sublayer supports fields in link-layer frames that enable multiple higher layer protocols to share a single physical link.

The MAC sublayer governs protocol access to the physical network medium. Through the MAC addresses that are typically assigned to all ports on a device, multiple devices on the same physical link can uniquely identify one another at the Data Link Layer. MAC addresses are used in addition to the network addresses that are typically configured manually on ports within a network.

MAC Addressing

A MAC address is the serial number permanently stored in a device adapter to uniquely identify the device. MAC addresses operate at the Data Link Layer, while IP addresses operate at the Network Layer. The IP address of a device can change as the device is moved around a network to different IP subnets, but the MAC address remains the same, because it is physically tied to the device.

Within an IP network, devices match each MAC address to its corresponding configured IP address by means of the Address Resolution Protocol (ARP). ARP maintains a table with a mapping for each MAC address in the network.

Most Layer 2 networks use one of three primary numbering spaces—MAC-48, EUI-48 (extended unique identifier), and EUI-64—which are all globally unique. MAC-48 and EUI-48 spaces each use 48-bit

addresses, and EUI-64 spaces use a 64-bit addresses, but all three use the same numbering format. MAC-48 addresses identify network hardware, and EUI-48 addresses identify other devices and software.

The Ethernet and ATM technologies supported on devices use the MAC-48 address space. IPv6 uses the EUI-64 address space.

MAC-48 addresses are the most commonly used MAC addresses in most networks. These addresses are 12-digit hexadecimal numbers (48 bits in length) that typically appear in one of the following formats:

- *MM:MM:MM:SS:SS:SS*
- *MM-MM-MM-SS-SS-SS*

The first three octets (*MM:MM:MM* or *MM-MM-MM*) are the ID number of the hardware manufacturer. Manufacturer ID numbers are assigned by the Institute of Electrical and Electronics Engineers (IEEE). The last three octets (*SS:SS:SS* or *SS-SS-SS*) make up the serial number for the device, which is assigned by the manufacturer. For example, an Ethernet interface card might have a MAC address of `00:05:85:c1:a6:a0`.

Release History Table

Release	Description
15.1X49-D100	Starting with Junos OS Release 15.1X49-D100, SRX320, SRX340, SRX345, and SRX550HM devices support the LTE interface. The dialer interface is used for initiating wireless WAN connections over LTE networks.

Physical Interface Properties

IN THIS SECTION

- Understanding Interface Physical Properties | **15**
- Understanding Bit Error Rate Testing | **17**
- Understanding Interface Clocking | **18**
- Understanding Frame Check Sequences | **19**
- MTU Default and Maximum Values | **20**
- Understanding Jumbo Frames Support for Ethernet Interfaces | **24**

The physical interfaces on security devices affect the transmission of either link-layer signals or the data across the links. The topics below describes the physical properties that include clocking properties, transmission properties, such as the maximum transmission unit (MTU), and encapsulation methods, such as point-to-point and Frame Relay encapsulation. SRX series devices also support jumbo frames.

Understanding Interface Physical Properties

The physical properties of a network interface are the characteristics associated with the physical link that affect the transmission of either link-layer signals or the data across the links. Physical properties include clocking properties, transmission properties, such as the maximum transmission unit (MTU), and encapsulation methods, such as point-to-point and Frame Relay encapsulation.

The default property values for an interface are usually sufficient to successfully enable a bidirectional link. However, if you configure a set of physical properties on an interface, those same properties must be set on all adjacent interfaces to which a direct connection is made.

Table 6 on page 15 summarizes some key physical properties of device interfaces.

Table 6: Interface Physical Properties

Physical Property	Description
`bert-error-rate`	Bit error rate (BER). The error rate specifies the number of bit errors in a particular bit error rate test (BERT) period required to generate a BERT error condition. See "Understanding Bit Error Rate Testing" on page 17.
`bert-period`	Bit error rate test (BERT) time period over which bit errors are sampled. See "Understanding Bit Error Rate Testing" on page 17.
`chap`	Challenge Handshake Authentication Protocol (CHAP). Specifying `chap` enables CHAP authentication on the interface. See "Understanding CHAP Authentication on a PPPoE Interface" on page 328.
`clocking`	Clock source for the link. Clocking can be provided by the local system (internal) or a remote endpoint on the link (external). By default, all interfaces use the internal clocking mode. If an interface is configured to accept an external clock source, one adjacent interface must be configured to act as a clock source. Under this configuration, the interface operates in a loop timing mode, in which the clocking signal is unique for that individual network segment or loop. See "Understanding Interface Clocking" on page 18.

Table 6: Interface Physical Properties *(Continued)*

Physical Property	Description
`description`	A user-defined text description of the interface, often used to describe the interface's purpose.
`disable`	Administratively disables the interface.
`encapsulation`	Type of encapsulation on the interface. Common encapsulation types include PPP, Frame Relay, Cisco HDLC, and PPP over Ethernet (PPPoE). See "Understanding Physical Encapsulation on an Interface" on page 274.
`fcs`	Frame check sequence (FCS). FCS is an error-detection scheme that appends parity bits to a digital signal and uses decoding algorithms that detect errors in the received digital signal.
`mtu`	Maximum transmission unit (MTU) size. MTU is the largest size packet or frame, specified in bytes or octets, that can be sent in a packet-based or frame-based network. The TCP uses MTU to determine the maximum size of each packet in any transmission. You can adjust the MTU values at the physical interfaces by using the following command: `set interface` *`interface-name`* `mtu` *`mtu-value`* Sometimes there is a need to reduce the MTU values on interfaces to match the host tap interface MTU otherwise packets are dropped. You can adjust the MTU values by setting the `mtu` option of the `set interfaces [fxp0 | em0 | fab0 | fab1]` command to a value between 256 and 9192. Example: `user@host#` **`set interfaces em0 mtu 1400`** The supported range for configuring an MTU packet size is 256 through 9192 bytes. However, all interfaces do not support 9192 bytes. For more information on the supported interfaces, see "MTU Default and Maximum Values" on page 20.

Table 6: Interface Physical Properties *(Continued)*

Physical Property	Description
`no-keepalives`	Disabling of keepalive messages across a physical link. A keepalive message is sent between network devices to indicate that they are still active. Keepalives help determine whether the interface is operating correctly. Except for ATM-over-ADSL interfaces, all interfaces use keepalives by default.
`pap`	Password Authentication Protocol (PAP). Specifying `pap` enables PAP authentication on the interface. See "Understanding CHAP Authentication on a PPPoE Interface" on page 328.
`payload-scrambler`	Scrambling of traffic transmitted out the interface. Payload scrambling randomizes the data payload of transmitted packets. Scrambling eliminates nonvariable bit patterns (strings of all 1s or all 0s) that generate link-layer errors across some physical links.

Understanding Bit Error Rate Testing

In telecommunication transmission, the bit error rate (BER) is the percentage of bits that have errors compared to the total number of bits received in a transmission, usually expressed as 10 to a negative power. For example, a transmission with a BER of 10^{-6} received 1 errored bit in 1,000,000 bits transmitted. The BER indicates how often a packet or other data unit must be retransmitted because of an error. If the BER is too high, a slower data rate might improve the overall transmission time for a given amount of data if it reduces the BER and thereby lowers the number of resent packets.

A bit error rate test (BERT) is a procedure or device that measures the BER for a given transmission. You can configure a device to act as a BERT device by configuring the interface with a bit error rate and a testing period. When the interface receives a BERT request from a BER tester, it generates a response in a well-known BERT pattern. The initiating device checks the BERT-patterned response to determine the number of bit errors.

Understanding Interface Clocking

IN THIS SECTION

- Data Stream Clocking | 19
- Explicit Clocking Signal Transmission | 19

Clocking determines how individual routing nodes or entire networks sample transmitted data. As streams of information are received by a device in a network, a clock source specifies when to sample the data. In asynchronous networks, the clock source is derived locally, and synchronous networks use a central, external clock source. Interface clocking indicates whether the device uses asynchronous or synchronous clocking.

NOTE: Because truly synchronous networks are difficult to design and maintain, most synchronous networks are really plesiochronous networks. In a plesiochronous network, different timing regions are controlled by local clocks that are synchronized (with very narrow constraints). Such networks approach synchronicity and are generally known as synchronous networks.

Most networks are designed to operate as asynchronous networks. Each device generates its own clock signal, or devices use clocks from more than one clock source. The clocks within the network are not synchronized to a single clock source. By default, devices generate their own clock signals to send and receive traffic.

The system clock allows the device to sample (or detect) and transmit data being received and transmitted through its interfaces. Clocking enables the device to detect and transmit the 0s and 1s that make up digital traffic through the interface. Failure to detect the bits within a data flow results in dropped traffic.

Short-term fluctuations in the clock signal are known as . Long-term variations in the signal are known as .

Asynchronous clocking can either derive the clock signal from the data stream or transmit the clocking signal explicitly.

This topic contains the following sections:

Data Stream Clocking

Common in T1 links, data stream clocking occurs when separate clock signals are not transmitted within the network. Instead, devices must extract the clock signal from the data stream. As bits are transmitted across the network, each bit has a time slot of 648 nanoseconds. Within a time slot, pulses are transmitted with alternating voltage peaks and drops. The receiving device uses the period of alternating voltages to determine the clock rate for the data stream.

Explicit Clocking Signal Transmission

Clock signals that are shared by hosts across a data link must be transmitted by one or both endpoints on the link. In a serial connection, for example, one host operates as a clock primary and the other operates as a clock client. The clock primary internally generates a clock signal that is transmitted across the data link. The clock client receives the clock signal and uses its period to determine when to sample data and how to transmit data across the link.

This type of clock signal controls only the connection on which it is active and is not visible to the rest of the network. An explicit clock signal does not control how other devices or even other interfaces on the same device sample or transmit data.

Understanding Frame Check Sequences

IN THIS SECTION

- Cyclic Redundancy Checks and Checksums | **20**
- Two-Dimensional Parity | **20**

All packets or frames within a network can be damaged by crosstalk or interference in the network's physical wires. The frame check sequence (FCS) is an extra field in each transmitted frame that can be analyzed to determine if errors have occurred. The FCS uses cyclic redundancy checks (CRCs), checksums, and two-dimensional parity bits to detect errors in the transmitted frames.

This topic contains the following sections:

Cyclic Redundancy Checks and Checksums

On a link that uses CRCs for frame checking, the data source uses a predefined polynomial algorithm to calculate a CRC number from the data it is transmitting. The result is included in the FCS field of the frame and transmitted with the data. On the receiving end, the destination host performs the same calculation on the data it receives.

If the result of the second calculation matches the contents of the FCS field, the packet was sent and received without bit errors. If the values do not match, an FCS error is generated, the frame is discarded and the originating host is notified of the error.

Checksums function similarly to CRCs, but use a different algorithm.

Two-Dimensional Parity

On a link that uses two-dimensional parity bits for frame checking, the sending and receiving hosts examine each frame in the total packet transmission and create a parity byte that is evaluated to detect transmission errors.

For example, a host can create the parity byte for the following frame sequence by summing up each column (each bit position in the frame) and keeping only the least-significant bit:

```
Frame 1       0 1 0 1 0 0 1
Frame 2       1 1 0 1 0 0 1
Frame 3       1 0 1 1 1 1 0
Frame 4       0 0 0 1 1 1 0
Frame 5       0 1 1 0 1 0 0
Frame 6       1 0 1 1 1 1 1

Parity Byte   1 1 1 1 0 1 1
```

If the sum of the bit values in a bit position is even, the parity bit for the position is 0. If the sum is odd, the parity bit is 1. This method is called even parity. Matching parity bytes on the originating and receiving hosts indicate that the packet was received without error.

MTU Default and Maximum Values

The MTU values are by default without any MTU configurations. If the MTU value is set, then the formula `IFF MTU (IP MTU) = IFD MTU (Media MTU) - L2 Overhead` is applicable. See Table 7 on page 21 for default MTU values.

NOTE: For ATM MLPPP irrespective of UIFD MTU, the IP MTU is always 1500 because the IP MTU calculation is based on the LSQ interface. Even if you configure the LSQ family MTU, the IP MTU value cannot exceed 1504.

Table 7 on page 21 lists MTU values for the SRX Series Services Gateways Physical Interface Modules (PIMs).

Table 7: MTU Values for the SRX Series Services Gateways PIMs

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
1-Port Gigabit Ethernet Small Form-Factor Pluggable (SFP) Mini-PIM	1514	9010	1500
1-Port Small Form-Factor Pluggable (SFP) Mini-PIM	1514	1518	1500
DOCSIS Mini-PIM	1504	1504	1500
Serial Mini-PIM	1504	2000	1500
T1/E1 Mini-PIM	1504	2000	1500
Dual CT1/E1 GPIM	1504	9000	1500
Quad CT1/E1 GPIM	1504	9000	1500
2-Port 10- Gigabit Ethernet XPIM	1514	9192	1500
16-Port Gigabit Ethernet XPIM	1514	9192	1500

Table 7: MTU Values for the SRX Series Services Gateways PIMs *(Continued)*

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
24-Port Gigabit Ethernet XPIM	1514	9192	1500
ADSL2+ Mini-PIM (Encapsulation)			
`atm-snap`	1512	1512	1504
`atm-vcmux`	1512	1512	1512
`atm-nlpid`	1512	1512	1508
`atm-cisco-nlpid`	1512	1512	1510
`ether-over-atm-llc`	1512	1512	1488
`atm-ppp-llc`	1512	1512	1506
`atm-ppp-vcmux`	1512	1512	1510
`atm-mlppp-llc`	1512	1512	1500
`ppp-over-ether-over-atm-llc`	1512	1512	1480
VDSL- Mini-PIM AT mode (Encapsulation)			
`atm-snap`	1514	1514	1506

Table 7: MTU Values for the SRX Series Services Gateways PIMs *(Continued)*

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
atm-vcmux	1514	1514	1514
atm-nlpid	1514	1514	1510
atm-cisco-nlpid	1514	1514	1512
ether-over-atm-llc	1514	1524	1490
atm-ppp-llc	1514	1514	1508
atm-ppp-vcmux	1514	1514	1512
atm-mlppp-llc	1514	1514	1500
ppp-over-ether-over-atm-llc	1514	1514	1482
VDSL- Mini-PIM PT mode	1514	1514	1500
G.SHDSL Mini-PIM AT mode (Encapsulation)			
atm-snap	4482	4482	4470
atm-vcmux	4482	4482	4470
atm-nlpid	4482	4482	4470

Table 7: MTU Values for the SRX Series Services Gateways PIMs *(Continued)*

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
atm-cisco-nlpid	4482	4482	4470
ether-over-atm-llc	4482	4482	1500
atm-ppp-llc	4482	4482	4476
atm-ppp-vcmux	4482	4482	4480
atm-mlppp-llc	4482	4482	1500
ppp-over-ether-over-atm-llc	4482	4482	1492
G.SHDSL Mini-PIM PT mode	1514	1514	1500

Understanding Jumbo Frames Support for Ethernet Interfaces

SRX Series devices support jumbo frames up to 9192 bytes.

Jumbo frames are Ethernet frames with more than 1500 bytes of payload (maximum transmission unit [MTU]). Jumbo frames can carry up to 9000 bytes of payload.

You configure jumbo frames at the physical interface by using the following command:

`set interface *interface-name* mtu *mtu-value*`

Example:

```
user@host# set interfaces ge-0/0/0 mtu 9192
```

The supported range for configuring an MTU packet size is 256 through 9192 bytes. However, all interfaces do not support 9192 bytes. For more information on the supported interfaces, see "MTU Default and Maximum Values" on page 20.

Logical Interface Properties

IN THIS SECTION

- Understanding Interface Logical Properties | 25
- Understanding Protocol Families | 26

The logical interfaces can be configured on the security devices and the description is displayed in the output of the show commands. The logical properties of the security devices include protocol families, IP address or addresses associated with the interface, Virtual LAN (VLAN) tagging, and any firewall filters or routing policies.

Understanding Interface Logical Properties

The logical properties of an interface are the characteristics that do not apply to the physical interface or the wires connected to it. Logical properties include:

- Protocol families running on the interface (including any protocol-specific MTUs)
- IP address or addresses associated with the interface. A *logical interface* can be configured with an IPv6 address, IPv4 address, or both. The IP specification requires a unique address on every interface of each system attached to an IP network, so that traffic can be correctly routed. Individual hosts such as home computers must have a single IP address assigned. Devices must have a unique IP address for every interface.
- Virtual LAN (VLAN) tagging
- Any firewall filters or routing policies that are operating on the interface

SEE ALSO

Understanding Virtual LANs | 38

Understanding Protocol Families

IN THIS SECTION

Common Protocol Suites | 26
Other Protocol Suites | 27

A protocol family is a group of logical properties within an interface configuration. Protocol families include all the protocols that make up a protocol suite. To use a protocol within a particular suite, you must configure the entire protocol family as a logical property for an interface. The protocol families include common and not-so-common protocol suites.

This topic contains the following sections:

Common Protocol Suites

Junos OS protocol families include the following common protocol suites:

- Inet—Supports IP protocol traffic, including OSPF, BGP, and Internet Control Message Protocol (ICMP).
- Inet6—Supports IPv6 protocol traffic, including RIP for IPv6 (RIPng), IS-IS, and BGP.
- ISO—Supports IS-IS traffic.
- MPLS—Supports MPLS.

NOTE: Junos OS security features are flow-based—meaning the device sets up a flow to examine the traffic. Flow-based processing is not supported for ISO or MPLS protocol families.

Other Protocol Suites

In addition to the common protocol suites, Junos protocol families sometimes use the following protocol suites:

- `ccc`—Circuit cross-connect (CCC).
- `mlfr-uni-nni`—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI).
- `mlfr-end-to-end`—Multilink Frame Relay end-to-end.
- `mlppp`—Multilink Point-to-Point Protocol.
- `tcc`—Translational cross-connect (TCC).
- `tnp`—Trivial Network Protocol. This Juniper Networks proprietary protocol provides communication between the Routing Engine and the device's packet forwarding components. Junos OS automatically configures this protocol family on the device's internal interfaces only.

Understanding IPv4 and IPv6 Protocol Family

IN THIS SECTION

- Understanding IPv4 Addressing | **28**
- Understanding IPv6 Address Space, Addressing, Address Format, and Address Types | **32**
- Configuring the inet6 IPv6 Protocol Family | **36**

IPv4 addresses are 32-bit numbers that are typically displayed in dotted decimal notation and contains two primary parts: the network prefix and the host number. The topics below describes the IPv4 Classful Addressing, IPv4 Dotted Decimal Notation, IPv4 Subnetting, IPv4 Variable-Length Subnet Masks, understanding IP Version 6, IPv6 address types and use of them in Junos OS RX Series Services Gateway, and configuration of inet6 IPv6 Protocol Family.

Understanding IPv4 Addressing

IN THIS SECTION

- IPv4 Classful Addressing | 28
- IPv4 Dotted Decimal Notation | 29
- IPv4 Subnetting | 29
- IPv4 Variable-Length Subnet Masks | 31

IPv4 addresses are 32-bit numbers that are typically displayed in dotted decimal notation. A 32-bit address contains two primary parts: the network prefix and the host number.

All hosts within a single network share the same network address. Each host also has an address that uniquely identifies it. Depending on the scope of the network and the type of device, the address is either globally or locally unique. Devices that are visible to users outside the network (webservers, for example) must have a globally unique IP address. Devices that are visible only within the network must have locally unique IP addresses.

IP addresses are assigned by a central numbering authority called the Internet Assigned Numbers Authority (IANA). IANA ensures that addresses are globally unique where needed and has a large address space reserved for use by devices not visible outside their own networks.

This topic contains the following sections:

IPv4 Classful Addressing

To provide flexibility in the number of addresses distributed to networks of different sizes, 4-octet (32-bit) IP addresses were originally divided into three different categories or classes: class A, class B, and class C. Each address class specifies a different number of bits for its network prefix and host number:

- Class A addresses use only the first byte (octet) to specify the network prefix, leaving 3 bytes to define individual host numbers.
- Class B addresses use the first 2 bytes to specify the network prefix, leaving 2 bytes to define host addresses.
- Class C addresses use the first 3 bytes to specify the network prefix, leaving only the last byte to identify hosts.

In binary format, with an x representing each bit in the host number, the three address classes can be represented as follows:

```
00000000 xxxxxxxx xxxxxxxx xxxxxxxx (Class A)
00000000 00000000 xxxxxxxx xxxxxxxx (Class B)
00000000 00000000 00000000 xxxxxxxx (Class C)
```

Because each bit (x) in a host number can have a 0 or 1 value, each represents a power of 2. For example, if only 3 bits are available for specifying the host number, only the following host numbers are possible:

```
111 110 101 100 011 010 001 000
```

In each IP address class, the number of host-number bits raised to the power of 2 indicates how many host numbers can be created for a particular network prefix. Class A addresses have 2^{24} (or 16,777,216) possible host numbers, class B addresses have 2^{16} (or 65,536) host numbers, and class C addresses have 2^{8} (or 256) possible host numbers.

IPv4 Dotted Decimal Notation

The 32-bit IPv4 addresses are most often expressed in dotted decimal notation, in which each octet (or byte) is treated as a separate number. Within an octet, the rightmost bit represents 2^{0} (or 1), increasing to the left until the first bit in the octet is 2^{7} (or 128). Following are IP addresses in binary format and their dotted decimal equivalents:

```
11010000 01100010 11000000 10101010 = 208.98.192.170
01110110 00001111 11110000 01010101 = 118.15.240.85
00110011 11001100 00111100 00111011 = 51.204.60.59
```

IPv4 Subnetting

Because of the physical and architectural limitations on the size of networks, you often must break large networks into smaller subnetworks. Within a such a subnetted network, each interface requires its own network number and identifying subnet address.

NOTE: The IP routing world has shifted to Classless Inter-Domain Routing (CIDR). As its name implies, CIDR eliminates the notion of address classes and simply conveys a network prefix along

with a mask. The mask indicates which bits in the address identify the network (the prefix). This document discusses subnetting in the traditional context of classfull IP addresses.

Figure 1 on page 30 shows a network comprised of three subnets.

Figure 1: Subnets in a Network

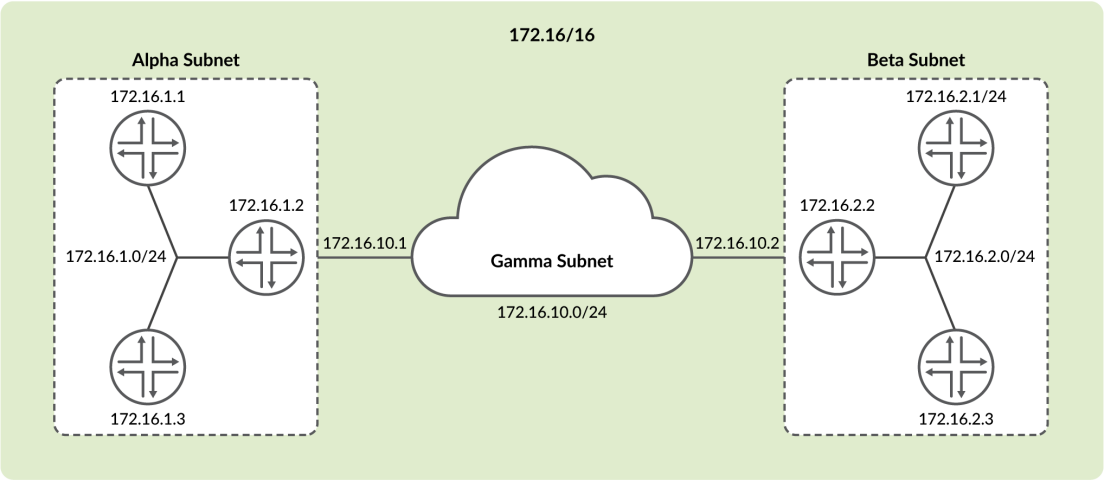

Figure 1 on page 30 shows three devices connected to the Alpha subnet on the left, three devices connected to the Beta subnet on the right, and a third subnet named Gamma that interconnects the left and right subnets over a WAN link. Collectively, the six devices and three subnets are contained within the larger class B network prefix. In this example, the organization is assigned the network prefix `172.16/16`, which is a class B address. Each subnet is assigned an IP address that falls within this class B network prefix.

In addition to sharing the class B network prefix (the first two octets), each subnet shares the third octet. Because we are using a /24 network mask in conjunction with a class B address, the third octet identifies the subnet. All devices on a subnet must have the same subnet address. In this case, the alpha subnet has the IP address `172.16.1.0/24`, the beta subnet has the IP address `172.16.2.0/24`, and the Gamma subnet is assigned `172.16.10.10/24`.

Taking one of these subnets as an example, the Beta subnet address `172.16.2.0/24` is represented in binary notation as:

```
10101100 . 00010000 . 00000010 . xxxxxxxx
```

Because the first 24 bits in the 32-bit address identify the subnet, the last 8 bits are available to assign to hosts attachments on each subnet. To reference a subnet, the address is written as `172.16.10.0/24` (or just `172.16.10/24`). The `/24` indicates the length of the subnet mask (sometimes written as `255.255.255.0`). This network mask indicates that the first 24 bits identify the network and subnetwork while the last 8 bits identify hosts on the respective subnetwork.

IPv4 Variable-Length Subnet Masks

Traditionally, subnets were divided by address class. Subnets had either 8, 16, or 24 significant bits, corresponding to 2^{24}, 2^{16}, or 2^{8} possible hosts. As a result, an entire /16 subnet had to be allocated for a network that required only 400 addresses, wasting 65,136 (2^{16} – 400 = 65,136) addresses.

To help allocate address spaces more efficiently, variable-length subnet masks (VLSMs) were introduced. Using VLSM, network architects can allocate more precisely the number of addresses required for a particular subnet.

For example, suppose a network with the prefix `192.14.17/24` is divided into two smaller subnets, one consisting of 18 devices and the other of 46 devices.

To accommodate 18 devices, the first subnet must have 2^{5} (32) host numbers. Having 5 bits assigned to the host number leaves 27 bits of the 32-bit address for the subnet. The IP address of the first subnet is therefore `192.14.17.128/27`, or the following in binary notation:

```
11000000 . 00001110 . 00010001 . 100xxxxx
```

The subnet mask includes 27 significant digits.

To create the second subnet of 46 devices, the network must accommodate 2^{6} (64) host numbers. The IP address of the second subnet is `192.14.17.64/26`, or

```
11000000 . 00001110 . 00010001 . 01xxxxxx
```

By assigning address bits within the larger `/24` subnet mask, you create two smaller subnets that use the allocated address space more efficiently.

Understanding IPv6 Address Space, Addressing, Address Format, and Address Types

IN THIS SECTION

- Understanding IP Version 6 (IPv6) | 32
- Understanding IPv6 Address Types and How Junos OS for SRX Series Services Gateway Uses Them | 33
- IPv6 Address Scope | 34
- IPv6 Address Structure | 34
- Understanding IPv6 Address Space, Addressing, and Address Types | 34
- Understanding IPv6 Address Format | 35

Understanding IP Version 6 (IPv6)

The ongoing expansive growth of the Internet and the need to provide IP addresses to accommodate it —to support increasing numbers of new users, computer networks, Internet-enabled devices, and new and improved applications for collaboration and communication—is escalating the emergent use of a new IP protocol. IPv6, with its robust architecture, was designed to satisfy these current and anticipated near future requirements.

IP version 4 (IPv4) is widely used throughout the world today for the Internet, intranets, and private networks. IPv6 builds upon the functionality and structure of IPv4 in the following ways:

- Provides a simplified and enhanced packet header to allow for more efficient routing.
- Improves support for mobile phones and other mobile computing devices.
- Enforces increased, mandatory data security through IPsec (which was originally designed for it).
- Provides more extensive quality-of-service (QoS) support.

IPv6 addresses consist of 128 bits, instead of 32 bits, and include a scope field that identifies the type of application suitable for the address. IPv6 does not support broadcast addresses, but instead uses multicast addresses for broadcast. In addition, IPv6 defines a new type of address called anycast.

Understanding IPv6 Address Types and How Junos OS for SRX Series Services Gateway Uses Them

IP version 6 (IPv6) includes the following types of addresses:

- Unicast

  A unicast address specifies an identifier for a single interface to which packets are delivered. Under IPv6, the vast majority of Internet traffic is foreseen to be unicast, and it is for this reason that the largest assigned block of the IPv6 address space is dedicated to unicast addressing. Unicast addresses include all addresses other than loopback, multicast, link-local-unicast, and unspecified.

  For SRX Series devices, the flow module supports the following kinds of IPv6 unicast packets:

  - Pass-through unicast traffic, including traffic from and to virtual routers. The device transmits pass-through traffic according to its routing table.
  - Host-inbound traffic from and to devices directly connected to SRX Series interfaces. For example, host-inbound traffic includes logging, routing protocol, and management types of traffic. The flow module sends these unicast packets to the Routing Engine and receives them from it. Traffic is processed by the Routing Engine instead of by the flow module, based on routing protocols defined for the Routing Engine.

    The flow module supports all routing and management protocols that run on the Routing Engine. Some examples are OSPFv3, RIPng, TELNET, and SSH.

- Multicast

  A multicast address specifies an identifier for a set of interfaces that typically belong to different nodes. It is identified by a value of 0xFF. IPv6 multicast addresses are distinguished from unicast addresses by the value of the high-order octet of the addresses.

  The devices support only host-inbound and host-outbound multicast traffic. Host inbound traffic includes logging, routing protocols, management traffic, and so on.

- Anycast

  An anycast address specifies an identifier for a set of interfaces that typically belong to different nodes. A packet with an anycast address is delivered to the nearest node, according to routing protocol rules.

  There is no difference between anycast addresses and unicast addresses except for the subnet-router address. For an anycast subnet-router address, the low order bits, typically 64 or more, are zero. Anycast addresses are taken from the unicast address space.

The flow module treats anycast packets in the same way as it handles unicast packets. If an anycast packet is intended for the device, it is treated as host-inbound traffic, and it delivers it to the protocol stack which continues processing it.

IPv6 Address Scope

Unicast and multicast IPv6 addresses support address scoping, which identifies the application suitable for the address.

Unicast addresses support global address scope and two types of local address scope:

- Link-local unicast addresses—Used only on a single network link. The first 10 bits of the prefix identify the address as a link-local address. Link-local addresses cannot be used outside the link.
- Site-local unicast addresses—Used only within a site or intranet. A site consists of multiple network links. Site-local addresses identify nodes inside the intranet and cannot be used outside the site.

Multicast addresses support 16 different types of address scope, including node, link, site, organization, and global scope. A 4-bit field in the prefix identifies the address scope.

IPv6 Address Structure

Unicast addresses identify a single interface. Each unicast address consists of n bits for the prefix, and 128 – n bits for the interface ID.

Multicast addresses identify a set of interfaces. Each multicast address consists of the first 8 bits of all 1s, a 4-bit flags field, a 4-bit scope field, and a 112-bit group ID:

```
11111111 | flgs | scop | group ID
```

The first octet of 1s identifies the address as a multicast address. The flags field identifies whether the multicast address is a well-known address or a transient multicast address. The scope field identifies the scope of the multicast address. The 112-bit group ID identifies the multicast group.

Similar to multicast addresses, anycast addresses identify a set of interfaces. However, packets are sent to only one of the interfaces, not to all interfaces. Anycast addresses are allocated from the normal unicast address space and cannot be distinguished from a unicast address in format. Therefore, each member of an anycast group must be configured to recognize certain addresses as anycast addresses.

Understanding IPv6 Address Space, Addressing, and Address Types

Addressing is the area where most of the differences between IP version 4 (IPv4) and IPv6 exist, but the changes are largely about the ways in which addresses are implemented and used. IPv6 has a vastly

larger address space than the impending exhausted IPv4 address space. IPv6 increases the size of the IP address from the 32 bits that compose an IPv4 address to 128 bits. Each extra bit given to an address doubles the size of the address space.

IPv4 has been extended using techniques such as Network Address Translation (NAT), which allows for ranges of private addresses to be represented by a single public address, and temporary address assignment. Although useful, these techniques fall short of the requirements of novel applications and environments such as emerging wireless technologies, always-on environments, and Internet-based consumer appliances.

In addition to the increased address space, IPv6 addresses differ from IPv4 addresses in the following ways:

- Includes a scope field that identifies the type of application that the address pertains to
- Does not support broadcast addresses, but instead uses multicast addresses to broadcast a packet
- Defines a new type of address, called anycast

Understanding IPv6 Address Format

All IPv6 addresses are 128 bits long, written as 8 sections of 16 bits each. They are expressed in hexadecimal representation, so the sections range from 0 to FFFF. Sections are delimited by colons, and leading zeroes in each section may be omitted. If two or more consecutive sections have all zeroes, they can be collapsed to a double colon.

IPv6 addresses consist of 8 groups of 16-bit hexadecimal values separated by colons (:). IPv6 addresses have the following format:

```
aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa
```

Each `aaaa` is a 16-bit hexadecimal value, and each `a` is a 4-bit hexadecimal value. Following is a sample IPv6 address:

```
3FFE:0000:0000:0001:0200:F8FF:FE75:50DF
```

You can omit the leading zeros of each 16-bit group, as follows:

```
3FFE:0:0:1:200:F8FF:FE75:50DF
```

You can compress 16-bit groups of zeros to double colons (::) as shown in the following example, but only once per address:

```
3FFE::1:200:F8FF:FE75:50DF
```

An IPv6 address prefix is a combination of an IPv6 prefix (address) and a prefix length. The prefix takes the form *ipv6-prefix*/*prefix-length* and represents a block of address space (or a network). The *ipv6-prefix* variable follows general IPv6 addressing rules. The *prefix-length* variable is a decimal value that indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address. For example, 10FA:6604:8136:6502::/64 is a possible IPv6 prefix with zeros compressed. The site prefix of the IPv6 address 10FA:6604:8136:6502::/64 is contained in the left most 64 bits, 10FA:6604:8136:6502.

For more information on the text representation of IPv6 addresses and address prefixes, see RFC 4291, *IP Version 6 Addressing Architecture*.

Limitations

SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices have the following limitations:

- Changes in source AS and destination AS are not immediately reflected in exported flows.
- IPv6 traffic transiting over IPv4 based IP over IP tunnel (for example, IPv6-over-IPv4 using ip-x/x/x interface) is not supported.

SEE ALSO

> About the IPv6 Basic Packet Header
>
> *Understanding IPv6 Packet Header Extensions*

Configuring the inet6 IPv6 Protocol Family

In configuration commands, the protocol family for IPv6 is named `inet6`. In the configuration hierarchy, instances of `inet6` are parallel to instances of `inet`, the protocol family for IPv4. In general, you configure `inet6` settings and specify IPv6 addresses in parallel to `inet` settings and IPv4 addresses.

NOTE: On SRX Series devices, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

The following example shows the CLI commands you use to configure an IPv6 address for an interface:

```
[edit]
user@host# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.100.37.178/24;
        }
    }
}

[edit]
user@host# set interfaces ge-0/0/0 unit 0 family ?
Possible completions:
+ apply-groups         Groups from which to inherit configuration data
+ apply-groups-except  Don't inherit configuration data from these groups
> ccc                  Circuit cross-connect parameters
> ethernet-switching   Ethernet switching parameters
> inet                 IPv4 parameters
> inet6                IPv6 protocol parameters
> iso                  OSI ISO protocol parameters
> mpls                 MPLS protocol parameters
> tcc                  Translational cross-connect parameters
> vpls                 Virtual private LAN service parameters

[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet6 address 8d8d:8d01::1/64
user@host# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.100.37.178/24;
        }
        family inet6 {
            address 8d8d:8d01::1/64;
        }
```

```
    }
}
```

SEE ALSO

Enabling Flow-Based Processing for IPv6 Traffic

Configuring VLAN Tagging

IN THIS SECTION

- Understanding Virtual LANs | **38**
- VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices | **40**
- Configuring VLAN Tagging | **41**

Virtual LANs (VLANs) allow network architects to segment LANs into different broadcast domains based on logical groupings. The topic below describes the configuration of these tagged VLANs, VLAN IDs, and supported Ethernet interface types on SRX series devices.

Understanding Virtual LANs

A LAN is a single broadcast domain. When traffic is broadcast, all hosts within the LAN receive the broadcast traffic. A LAN is determined by the physical connectivity of devices within the domain.

Within a traditional LAN, hosts are connected by a hub or repeater that propagates any incoming traffic throughout the network. Each host and its connecting hubs or repeaters make up a LAN segment. LAN

segments are connected through switches and bridges to form the broadcast domain of the LAN. Figure 2 on page 39 shows a typical LAN topology.

Figure 2: Typical LAN

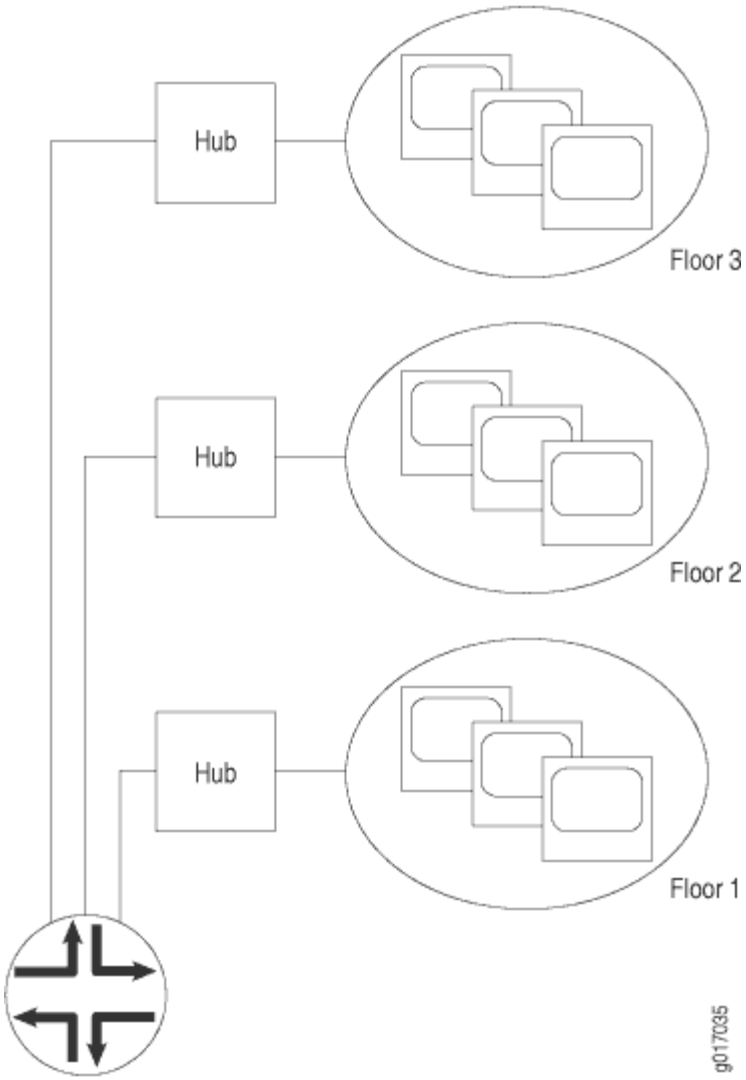

Virtual LANs (VLANs) allow network architects to segment LANs into different broadcast domains based on logical groupings. Because the groupings are logical, the broadcast domains are not determined by the physical connectivity of the devices in the network. Hosts can be grouped according to a logical function, to limit the traffic broadcast within the VLAN to only the devices for which the traffic is intended.

Suppose a corporate network has three major organizations: engineering, sales, and support. Using VLAN tagging, hosts within each organization can be tagged with a different VLAN identifier. Traffic sent

to the broadcast domain is then checked against the VLAN identifier and broadcast to only the devices in the appropriate VLAN. Figure 3 on page 40 shows a typical VLAN topology.

Figure 3: Typical VLAN

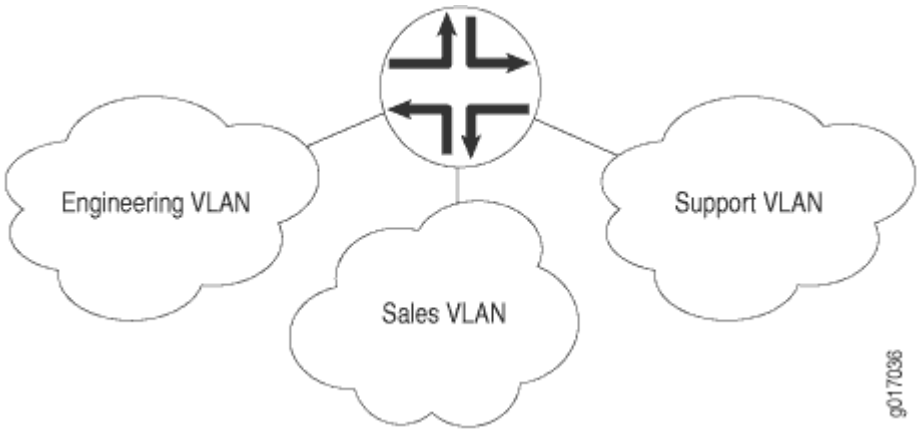

SEE ALSO

MPLS Applications User Guide

VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices

Table 8 on page 40 lists VLAN ID range by interface type supported on SRX Series devices:

Table 8: VLAN ID Range by Interface Type Supported on the SRX Series Devices

Interface Type	Interface Type VLAN ID Range
2-Port 10-Gigabit Ethernet	1 through 4094
10-Gigabit Ethernet	1 through 4094
16-Port Gigabit Ethernet	1 through 4094

Table 8: VLAN ID Range by Interface Type Supported on the SRX Series Devices *(Continued)*

Interface Type	Interface Type VLAN ID Range
24-Port Gigabit Ethernet	1 through 4094
Aggregated Ethernet for Fast Ethernet	1 through 1023
Aggregate Ethernet for Gigabit Ethernet	1 through 4094
Gigabit Ethernet	1 through 4094
Management and internal Ethernet interfaces	1 through 1023

NOTE: On SRX210, SRX220, SRX240, SRX320, and SRX340 devices, on 1-GE SFP Mini-PIM, the VLAN ID 4093 falls under the reserved VLAN address range. (Platform support depends on the Junos OS release in your installation.) Because of this, you will not be able to configure VLAN ID from this range.

SEE ALSO

Understanding Interface Physical Properties | 15

Configuring VLAN Tagging

IN THIS SECTION

- Configuring Single-Tag Framing | 42
- Configuring Dual Tagging | 42
- Configuring Mixed Tagging | 43
- Configuring Mixed Tagging Support for Untagged Packets | 44

You can configure SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices to receive and forward single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.

See Table 9 on page 42 for flexible VLANs.

Table 9: Flexible VLANs

Number of Tags	VLAN ID
0 (Untagged)	Native
1 (Tagged)	Single
2 (Dual tagged)	Dual

This topic includes the following sections:

Configuring Single-Tag Framing

To configure a device to receive and forward single-tag frames with 802.1Q VLAN tags, include the `vlan-tagging` statement at the `[edit interfaces` *`interface-name`*`]` hierarchy level:

```
[edit interfaces interface-name]
vlan-tagging;
```

NOTE: SRX5400, SRX5600, and SRX5800 only support single-tag framing.

Configuring Dual Tagging

To configure the device to receive and forward dual-tag frames with 802.1Q VLAN tags, include the `flexible-vlan-tagging` statement at the `[edit interfaces` *`interface-name`*`]` hierarchy level:

```
[edit interfaces interface-name]
flexible-vlan-tagging;
```

Configuring Mixed Tagging

Mixed tagging is supported on ethernet interfaces of SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices. Mixed tagging lets you configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.

To configure mixed tagging, include the `flexible-vlan-tagging` statement at the `[edit interfaces ge-`*`fpc/pic/port`*` ]` hierarchy level. You must also include the `vlan-tags` statement with `inner` and `outer` options or the `vlan-id` statement at the `[edit interfaces ge-`*`fpc/pic/port`*` unit `*`logical-unit-number`*`]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
unit logical-unit-number {
    vlan-id number;
    family family {
        address address;
    }
}
unit logical-unit-number {
    vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
    family family {
        address address;
    }
}
```

NOTE: When you configure the physical interface MTU for mixed tagging, you must increase the MTU to 4 bytes more than the MTU value you would configure for a standard VLAN-tagged interface.

For example, if the MTU value is configured to be 1018 on a VLAN-tagged interface, then the MTU value on a flexible VLAN tagged interface must be 1022—4 bytes more. The additional 4 bytes accommodates the future addition of a stacked VLAN tag configuration on the same physical interface.

The following example configures mixed tagging. Dual-tag and single-tag logical interfaces are under the same physical interface:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
unit 0 {
```

```
    vlan-id 232;
    family inet {
        address 10.66.1.2/30;
    }
}
unit 1 {
    vlan-tags outer 0x8100.222 inner 0x8100.221;
    family inet {
        address 10.66.1.2/30;
    }
}
```

Configuring Mixed Tagging Support for Untagged Packets

You can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untagged packets, include the `native-vlan-id` statement and the `flexible-vlan-tagging` statement at the `[edit interfaces` *`interface-name`*`]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
native-vlan-id number;
```

NOTE: The `flexible-vlan-tagging` is supported only with either no encapsulation or VPLS VLAN encapsulation.

The *logical interface* on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the `vlan-id` statement (matching the `native-vlan-id` statement on the physical interface) at the `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*`]` hierarchy level.

The following example configures untagged packets to be mapped to logical unit number 0:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
native-vlan-id 232;
unit 0 {
    vlan-id 232;
    family inet {
        address 10.66.1.2/30;
```

```
    }
}
unit 1 {
    vlan-tags outer 0x8100.222 inner 0x8100.221;
    family inet {
        address 10.66.1.2/30;
    }
}
```

CHAPTER 2

Configuring DS1, DS3, and 1-Port Clear Channel DS3/E3 GPIM Interfaces

Configuring DS1 Interfaces | 47

Configuring DS3 Interfaces | 57

Configuring 1-Port Clear Channel DS3/E3 GPIM | 68

Configuring DS1 Interfaces

IN THIS SECTION

- Understanding T1 and E1 Interfaces | 47
- Example: Configuring a T1 Interface | 51
- Example: Deleting a T1 Interface | 55

T1 and E1 refer to the data transmission formats that carry DS1 signals across interfaces. The below topic discuss the functionality of T1 and E1, configuration details and also deleting the T1 interface.

Understanding T1 and E1 Interfaces

IN THIS SECTION

- T1 Overview | 48
- E1 Overview | 48
- T1 and E1 Signals | 48
- Encoding | 49
- T1 and E1 Framing | 50
- T1 and E1 Loopback Signals | 50

T1 and E1 are equivalent digital data transmission formats that carry DS1 signals. T1 and E1 lines can be interconnected for international use.

This topic contains the following sections:

T1 Overview

T1 is a digital data transmission medium capable of handling 24 simultaneous connections running at a combined 1.544 Mbps. T1 combines these 24 separate connections, called channels or time slots, onto a single link. T1 is also called DS1.

The T1 data stream is broken into frames. Each frame consists of a single framing bit and 24 8-bit channels, totaling 192 bits per T1 frame. Frames are transmitted 8,000 times per second, at a data transmission rate of 1.544 Mbps (8,000 x 193 = 1.544 Mbps).

As each frame is received and processed, the data in each 8-bit channel is maintained with the channel data from previous frames, enabling T1 traffic to be separated into 24 separate flows across a single medium. For example, in the following set of 4-channel frames (without a framing bit), the data in channel 1 consists of the first octet of each frame, the data in channel 2 consists of the second octet of each frame, and so on:

```
           Chan. 1   Chan. 2   Chan. 3   Chan. 4
Frame 1  [10001100][00110001][11111000][10101010]
Frame 2  [11100101][01110110][10001000][11001010]
Frame 3  [00010100][00101111][11000001][00000001]
```

E1 Overview

E1 is the European format for DS1 digital transmission. E1 links are similar to T1 links except that they carry signals at 2.048 Mbps. Each signal has 32 channels, and each channel transmits at 64 Kbps. E1 links have higher bandwidth than T1 links because it does not reserve one bit for overhead. Whereas, T1 links use 1 bit in each channel for overhead.

T1 and E1 Signals

T1 and E1 interfaces consist of two pairs of wires—a transmit data pair and a receive data pair. Clock signals, which determine when the transmitted data is sampled, are embedded in the T1 and E1 transmissions.

Typical digital signals operate by sending either zeros (0s) or ones (1s), which are usually represented by the absence or presence of a voltage on the line. The receiving device need only detect the presence of the voltage on the line at the particular sampling edge to determine whether the signal is 0 or 1. T1 and E1, however, use bipolar electrical pulses. Signals are represented by no voltage (0), positive voltage (1), or negative voltage (1). The bipolar signal allows T1 and E1 receivers to detect error conditions in the line, depending on the type of encoding that is being used.

Encoding

The following are common T1 and E1 encoding techniques:

- Alternate mark inversion (AMI)—T1 and E1
- Bipolar with 8-zero substitution (B8ZS)—T1 only
- High-density bipolar 3 code (HDB3)—E1 only

AMI Encoding

AMI encoding forces the 1s signals on a T1 or E1 line to alternate between positive and negative voltages for each successive 1 transmission, as in this sample data transmission:

```
1 1 0 1 0 1 0 1
+ - 0 + 0 - 0 +
```

When AMI encoding is used, a data transmission with a long sequence of 0s has no voltage transitions on the line. In other words, voice transmission does not use AMI encoding because it never encounters the “long string of zeroes” problem. In this situation, devices have difficulty maintaining clock synchronization, because they rely on the voltage fluctuations to constantly synchronize with the transmitting clock. To counter this effect, the number of consecutive 0s in a data stream is restricted to 15. This restriction is called the 1s density requirement, because it requires a certain number of 1s for every 15 0s that are transmitted.

On an AMI-encoded line, two consecutive pulses of the same polarity—either positive or negative—are called a bipolar violation (BPV), which is generally flagged as an error.

B8ZS and HDB3 Encoding

Neither B8ZS nor HDB3 encoding restricts the number of 0s that can be transmitted on a line. Instead, these encoding methods detect sequences of 0s and substitute bit patterns for the sequences to provide the signal oscillations required to maintain timing on the link.

The B8ZS encoding method for T1 lines detects sequences of eight consecutive 0 transmissions and substitutes a pattern of two consecutive BPVs (11110000). Because the receiving end uses the same encoding, it detects the BPVs as 0s substitutions, and no BPV error is flagged. A single BPV, which does not match the 11110000 substitution bit sequence is likely to generate an error, depending on the configuration of the device.

B8ZS uses bipolar violations to synchronize devices, a solution that does not require the use of extra bits, which means a T1 circuit using B8ZS can use the full 64 Kbps for each channel for data.

The HDB3 encoding method for E1 lines detects sequences of four consecutive 0 transmissions and substitutes a single BPV (1100). Similar to B8ZS encoding, the receiving device detects the 0s substitutions and does not generate a BPV error.

T1 and E1 Framing

T1 interfaces uses extended superframe (ESF). E1 interfaces use G.704 framing or G.704 with no CRC4 framing, or can be in unframed mode.

ESF Framing for T1

ESF extends the D4 superframe from 12 frames to 24 frames. By expanding the size of the superframe, ESF increases the number of bits in the superframe framing pattern from 12 to 24. The extra bits are used for frame synchronization, error detection, and maintenance communications through the facilities data link (FDL).

The ESF pattern for synchronization bits is 001011. Only the framing bits from frames 4, 8, 12, 16, 20, and 24 in the superframe sequence are used to create the synchronization pattern.

The framing bits from frames 2, 6, 10, 14, 18, and 22 are used to pass a CRC code for each superframe block. The CRC code verifies the integrity of the received superframe and detects bit errors with a CRC6 algorithm.

The framing bits for frames 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, and 23 are used for the data link channel. These 12 bits enable the operators at the network control center to query the remote equipment for information about the performance of the link.

T1 and E1 Loopback Signals

The control signal on a T1 or E1 link is the loopback signal. Using the loopback signal, the operators at the network control center can force the device at the remote end of a link to retransmit its received signals back onto the transmit path. The transmitting device can then verify that the received signals match the transmitted signals, to perform end-to-end checking on the link.

Two loopback signals are used to perform the end-to-end testing:

- The loop-up command signal sets the link into loopback mode, with the following command pattern:

```
...100001000010000100...
```

- The loop-down signal returns the link to its normal mode, with the following command pattern:

```
...100100100100100100...
```

While the link is in loopback mode, the operator can insert test equipment onto the line to test its operation.

Example: Configuring a T1 Interface

IN THIS SECTION

- Requirements | 51
- Overview | 51
- Configuration | 52
- Verification | 53

This example shows how to complete the initial configuration on a T1 interface.

Requirements

Before you begin, install a PIM, connect the interface cables to the ports, and power on the device. See the *Getting Started Guide* for your device.

Overview

This example describes the initial configuration that you must complete on each network interface. In this example, you configure the t1-1/0/0 interface as follows:

- You create the basic configuration for the new interface by setting the encapsulation type to ppp. You can enter additional values for physical interface properties as needed.
- You set the logical interface to 0. Note that the logical unit number can range from 0 through 16,384. You can enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

IN THIS SECTION

- Procedure | 52

Procedure

CLI Quick Configuration

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

```
set interfaces t1-1/0/0 encapsulation ppp unit 0
```

Step-by-Step Procedure

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

To configure a T1 interface:

1. Create the interface.

```
[edit]
user@host# edit interfaces t1-1/0/0
```

2. Create the basic configuration for the new interface.

```
[edit interfaces t1-1/0/0]
user@host# set encapsulation ppp
```

3. Add logical interfaces.

```
[edit interfaces t1-1/0/0]
user@host# set unit 0
```

Results

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

For brevity, this `show interfaces` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
    ...
    t1-1/0/0 {
    encapsulation ppp;
    unit 0;
}
```

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

Verification

IN THIS SECTION

- Verifying the Link State of All Interfaces | 54
- Verifying Interface Properties | 54

Confirm that the configuration is working properly.

Verifying the Link State of All Interfaces

Purpose

By using the ping tool on each peer address in the network, verify that all interfaces on the device are operational.

Action

For each interface on the device:

1. In the J-Web interface, select `Troubleshoot>Ping Host`.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click `Start`. The output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time, in milliseconds, is listed in the time field.

Meaning

Verifying Interface Properties

Purpose

Verify that the interface properties are correct.

Action

From the operational mode, enter the `show interfaces detail` command.

The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do one of the following:

- In the CLI configuration editor, delete the `disable` statement at the [edit interfaces t1-1/0/0] level of the configuration hierarchy.
- In the J-Web configuration editor, clear the `Disable` check box on the Interfaces> t1-1/0/0 page.

- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. It indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of input and output bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the `clear interfaces statistics t1-1/0/0` command.

Example: Deleting a T1 Interface

IN THIS SECTION

- Requirements | 55
- Overview | 55
- Configuration | 56
- Verification | 56

This example shows how to delete a T1 interface.

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you delete the t1-1/0/0 interface.

NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on the J-Web pages.

Configuration

IN THIS SECTION

- Procedure | 56

Procedure

Step-by-Step Procedure

To delete a T1 interface:

1. Specify the interface you want to delete.

```
[edit interfaces]
user@host# delete t1-1/0/0
```

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces` command.

Configuring DS3 Interfaces

IN THIS SECTION

- Understanding T3 and E3 Interfaces | 57
- Example: Configuring a T3 Interface | 63
- Example: Deleting a T3 Interface | 67

DS3 interfaces, also referred to as T3, is an high-speed data transmission medium formed by multiplexing DS1 and DS2 signals. The below topic discuss the functionality of T3 interfaces, configuration details and also deleting the T3 interface.

Understanding T3 and E3 Interfaces

IN THIS SECTION

- Multiplexing DS1 Signals | 58
- DS2 Bit Stuffing | 58
- DS3 Framing | 59

T3 is a high-speed data-transmission medium formed by multiplexing 28 DS1 signals into seven separate DS2 signals, and combining the DS2 signals into a single DS3 signal. T3 links operate at 43.736 Mbps. T3 is also called DS3.

E3 is the equivalent European transmission format. E3 links are similar to T3 (DS3) links, but carry signals at 34.368 Mbps. Each signal has 16 E1 channels, and each channel transmits at 2.048 Mbps. E3 links use all 8 bits of a channel, whereas T3 links use 1 bit in each channel for overhead.

Multiplexing DS1 Signals

Four DS1 signals combine to form a single DS2 signal. The four DS1 signals form a single DS2 M-frame, which includes subframes M1 through M4. Each subframe has six 49-bit blocks, for a total of 294 bits per subframe. The first bit in each block is a DS2 overhead (OH) bit. The remaining 48 bits are DS1 information bits.

Figure 4 on page 58 shows the DS2 M-frame format.

Figure 4: DS2 M-Frame Format

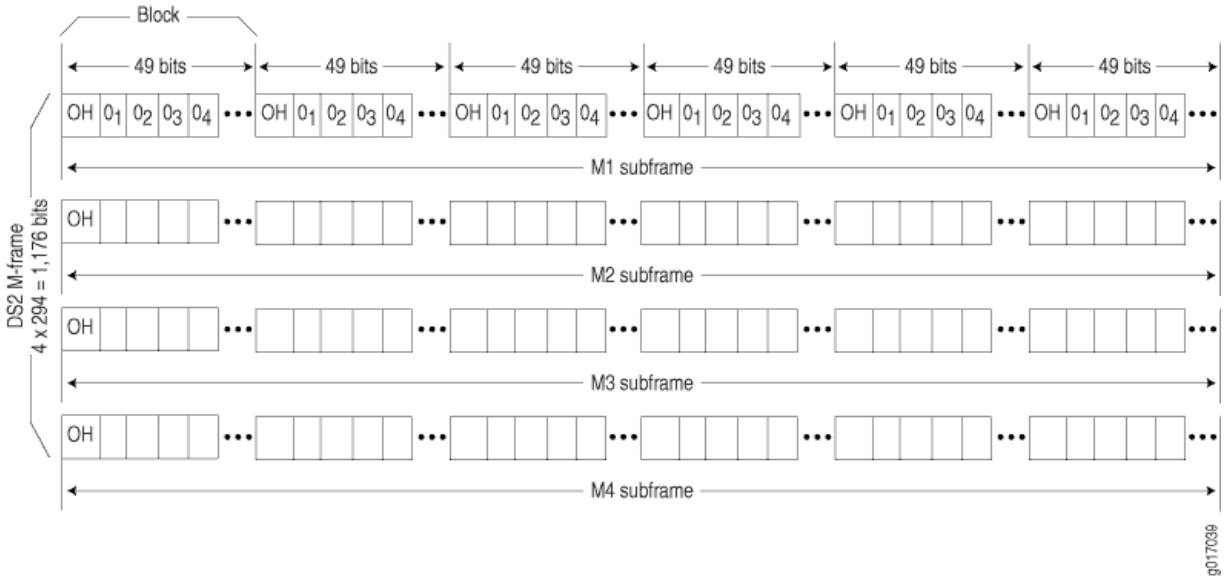

The four DS2 subframes are not four DS1 channels. Instead, the DS1 data bits within the subframes are formed by data interleaved from the DS1 channels. The 0*n* values designate time slots devoted to DS1 inputs as part of the bit-by-bit interleaving process. After every 48 DS1 information bits (12 bits from each signal), a DS2 OH bit is inserted to indicate the start of a subframe.

DS2 Bit Stuffing

Because the four DS1 signals are asynchronous signals, they might operate at different line rates. To synchronize the asynchronous streams, the multiplexers on the line use bit stuffing.

A DS2 connection requires a nominal transmit rate of 6.304 Mbps. However, because multiplexers increase the overall output rate to the intermediate rate of 6.312 Mbps, the output rate is higher than individual input rates on DS1 signals. The extra bandwidth is used to stuff the incoming DS1 signals with extra bits until the output rate of each signal equals the increased intermediate rate. These stuffed bits are inserted at fixed locations in the DS2 M-frame. When DS2 frames are received and the signal is demultiplexed, the stuffing bits are identified and removed.

DS3 Framing

A set of four DS1 signals is multiplexed into seven DS2 signals, which are multiplexed into a single DS3 signal. The multiplexing occurs just as with DS1-to-DS2 multiplexing. The resulting DS3 signal uses either the standard M13 asynchronous framing format or the C-bit parity framing format. Although the two framing formats differ in their use of control and message bits, the basic frame structures are identical. The DS3 frame structures are shown in Figure 5 on page 59 and Figure 6 on page 61.

M13 Asynchronous Framing

A DS3 M-frame includes seven subframes, formed by DS2 data bits interleaved from the seven multiplexed DS2 signals. Each subframe has eight 85-bit blocks—a DS3 OH bit plus 84 data bits. The meaning of an OH bit depends on the block it precedes. Standard DS3 M13 asynchronous framing format is shown in Figure 5 on page 59.

Figure 5: DS3 M13 Frame Format

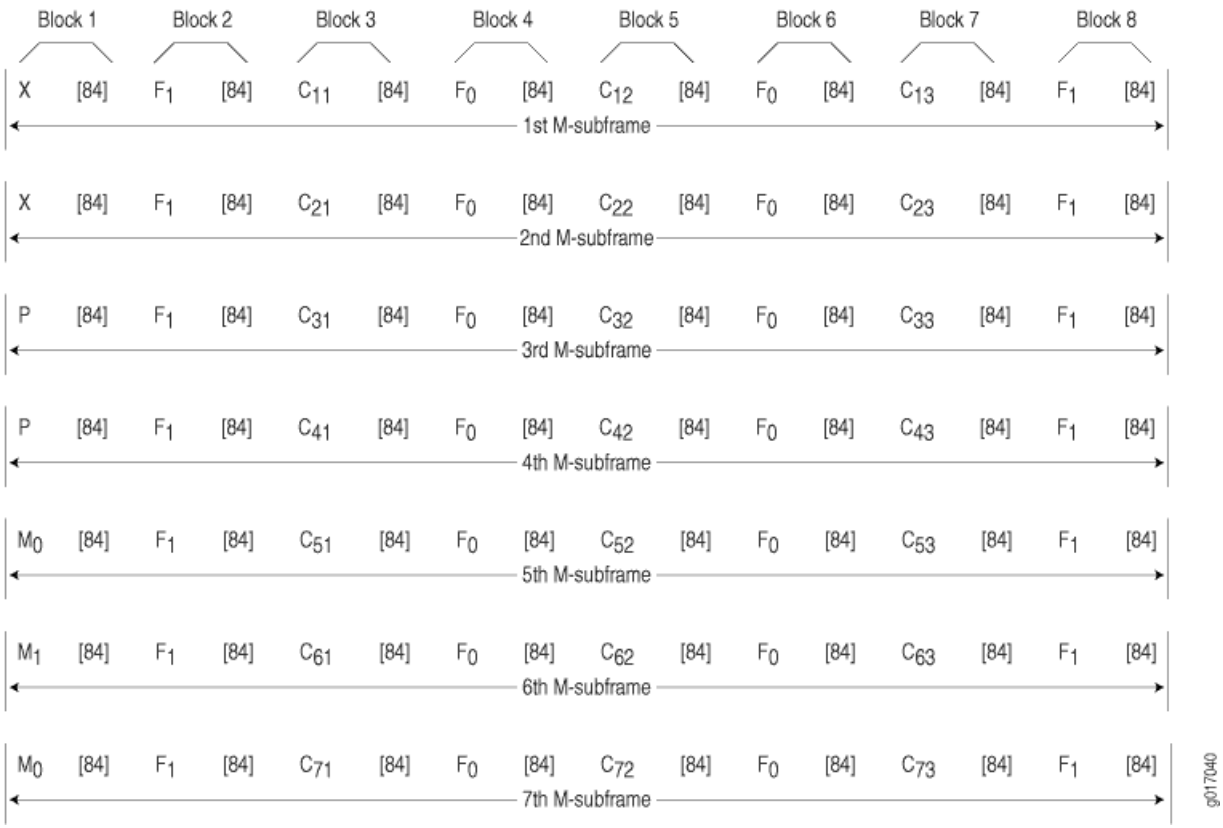

A DS3 M13 M-frame contains the following types of OH bits:

- Framing bits (F-bits)—Make up a frame alignment signal that synchronizes DS3 subframes. Each DS3 frame contains 28 F-bits (4 bits per subframe). F-bits are located at the beginning of blocks 2, 4, 6, and 8 of each subframe. When combined, the frame alignment pattern for each subframe is 1001. The pattern can be examined to detect bit errors in the transmission.
- Multiframing bits (M-bits)—Make up a multiframe alignment signal that synchronizes the M-frames in a DS3 signal. Each DS3 frame contains 3 M-bits, which are located at the beginning of subframes 5, 6, and 7. When combined, the multiframe alignment patter for each M-frame is 010.
- Bit stuffing control bits (C-bits)—Serve as bit stuffing indicators for each DS2 input. For example, C_{11}, C_{12}, and C_{13} are indicators for DS2 input 1. Their values indicate whether DS3 bit stuffing has occurred at the multiplexer. If the three C-bits in a subframe are all 0s, no stuffing was performed for the DS2 input. If the three C-bits are all 1s, stuffing was performed.
- Message bits (X-bits)—Used by DS3 transmitters to embed asynchronous in-service messages in the data transmission. Each DS3 frame contains 2 X-bits, which are located at the beginning of subframes 1 and 2. Within an DS3 M-frame, both X-bits must be identical.
- Parity bits (P-bits)—Compute parity over all but 1 bit of the M-frame. (The first X-bit is not included.) Each DS3 frame contains 2 P-bits, which are located at the beginning of subframes 3 and 4. Both P-bits must be identical.

  If the previous DS3 frame contained an odd number of 1s, both P-bits are set to 1. If the previous DS3 contained an even number of 1s, both P-bits are set to 0. If, on the receiving side, the number of 1s for a given frame does not match the P-bits in the following frame, it indicates one or more bit errors in the transmission.

C-Bit Parity Framing

In M13 framing, every C-bit in a DS3 frame is used for bit stuffing. However, because multiplexers first use bit stuffing when multiplexing DS1 signals into DS2 signals, the incoming DS2 signals are already synchronized. Therefore, the bit stuffing that occurs when DS2 signals are multiplexed is redundant.

C-bit parity framing format redefines the function of C-bits and X-bits, using them to monitor end-to-end path performance and provide in-band data links. The C-bit parity framing structure is shown in Figure 6 on page 61.

Figure 6: DS3 C-Bit Parity Framing

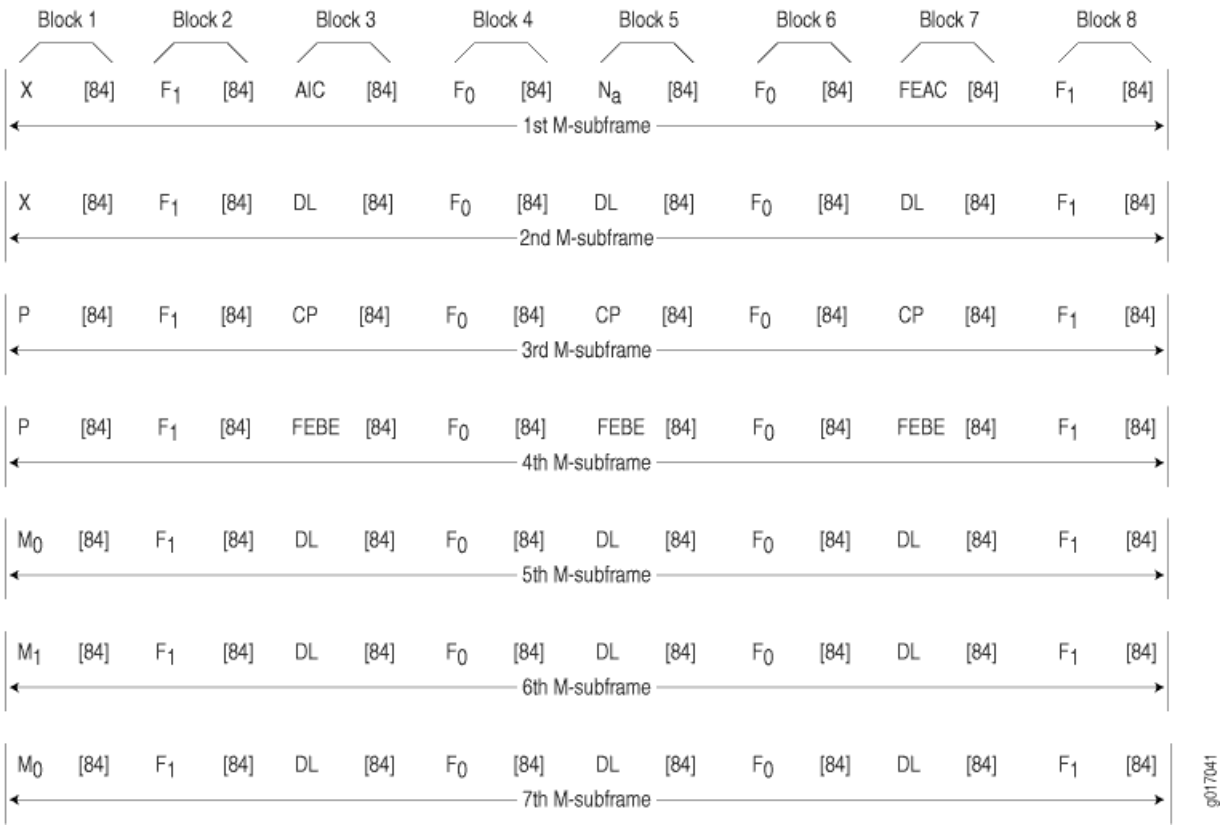

In C-bit parity framing, the X-bits transmit error conditions from the far end of the link to the near end. If no error conditions exist, both X-bits are set to 1. If an out-of-frame (OOF) or alarm indication signal (AIS) error is detected, both X-bits are set to 0 in the upstream direction for 1 second to notify the other end of the link about the condition.

The C-bits that control bit stuffing in M13 frames are typically used in the following ways by C-bit parity framing:

- Application identification channel (AIC)—The first C-bit in the first subframe identifies the type of DS3 framing used. A value of 1 indicates that C-bit parity framing is in use.
- N_a—A reserved network application bit.
- Far-end alarm and control (FEAC) channel—The third C-bit in the first subframe is used for the FEAC channel. In normal transmissions, the FEAC C-bit transmits all 1s. When an alarm condition is

present, the FEAC C-bit transmits a code word in the format `0xxxxxxx 11111111`, in which *x* can be either 1 or 0. Bits are transmitted from right to left.

Table 10 on page 62 lists some C-bit code words and the alarm or status condition indicated.

Table 10: FEAC C-Bit Condition Indicators

Alarm or Status Condition	C-Bit Code Word
DS3 equipment failure requires immediate attention.	`00110010 11111111`
DS3 equipment failure occurred—such as suspended, not activated, or unavailable service—that is non-service-affecting.	`00011110 11111111`
DS3 loss of signal.	`00011100 11111111`
DS3 out of frame.	`00000000 11111111`
DS3 alarm indication signal (AIS) received.	`00101100 11111111`
DS3 idle received.	`00110100 11111111`
Common equipment failure occurred that is non-service-affecting.	`00011101 11111111`
Multiple DS1 loss of signal.	`00101010 11111111`
DS1 equipment failure occurred that requires immediate attention.	`00001010 11111111`
DS1 equipment failure occurred that is non-service-affecting.	`00000110 11111111`
Single DS1 loss of signal.	`00111100 11111111`

- Data links—The 12 C-bits in subframes 2, 5, 6, and 7 are data link (DL) bits for applications and terminal-to-terminal path maintenance.
- DS3 parity—The 3 C-bits in the third subframe are DS3 parity C-bits (also called CP-bits). When a DS3 frame is transmitted, the sending device sets the CP-bits to the same value as the P-bits. When

the receiving device processes the frame, it calculates the parity of the M-frame and compares this value to the parity in the CP-bits of the following M-frame. If no bit errors have occurred, the two values are typically the same.

- Far–end block errors (FEBEs)—The 3 C-bits in the fourth subframe make up the far-end block error (FEBE) bits. If a framing or parity error is detected in an incoming M-frame (via the CP-bits), the receiving device generates a C-bit parity error and sends an error notification to the transmitting (far-end) device. If an error is generated, the FEBE bits are set to 000. If no error occurred, the bits are set to 111.

Example: Configuring a T3 Interface

IN THIS SECTION

- Requirements | 63
- Overview | 63
- Configuration | 64
- Verification | 65

This example shows how to complete the initial configuration on a T3 interface.

Requirements

Before you begin, install a PIM, connect the interface cables to the ports, and power on the device. See the *Getting Started Guide* for your device.

Overview

This example describes the initial configuration that you must complete on each network interface. In this example, you configure the t3-1/0/0 interface as follows:

- You create the basic configuration for the new interface by setting the encapsulation type to ppp. You can enter additional values for physical interface properties as needed.
- You set the logical interface to 0. Note that the logical unit number can range from 0 to 16,384. You can enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

IN THIS SECTION

- Procedure | 64

Procedure

CLI Quick Configuration

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

```
set interfaces t3-1/0/0 encapsulation ppp unit 0
```

Step-by-Step Procedure

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

To configure a T3 interface:

1. Create the interface.

```
[edit]
user@host# edit interfaces t3-1/0/0
```

2. Create the basic configuration for the new interface.

```
[edit interfaces t3-1/0/0]
user@host# set encapsulation ppp
```

3. Add logical interfaces.

```
[edit interfaces t3-1/0/0]
user@host# set unit 0
```

Results

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

For brevity, this `show interfaces` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
    ...
    t3-1/0/0 {
    encapsulation ppp;
    unit 0;
}
```

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

Verification

IN THIS SECTION

- Verifying the Link State of All Interfaces | 66
- Verifying Interface Properties | 66

Confirm that the configuration is working properly.

Verifying the Link State of All Interfaces

Purpose

By using the ping tool on each peer address in the network, verify that all interfaces on the device are operational.

Action

For each interface on the device:

1. In the J-Web interface, select `Troubleshoot>Ping Host`.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click `Start`. The output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time in milliseconds is listed in the time field.

Verifying Interface Properties

Purpose

Verify that the interface properties are correct.

Action

From the operational mode, enter the `show interfaces detail` command.

The output shows a summary of interface information. Verify the following information:

- The physical interface is Enabled. If the interface is shown as Disabled, do one of the following:
  - In the CLI configuration editor, delete the `disable` statement at the [edit interfaces t3-1/0/0] level of the configuration hierarchy.
  - In the J-Web configuration editor, clear the `Disable` check box on the Interfaces> t3-1/0/0 page.

- The physical link is Up. A link state of Down indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The Last Flapped time is an expected value. It indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of input and output bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the `clear interfaces statistics t3-1/0/0` command.

Example: Deleting a T3 Interface

IN THIS SECTION

- Requirements | 67
- Overview | 67
- Configuration | 68
- Verification | 68

This example shows how to delete a T3 interface.

Requirements

No special configuration beyond device initialization is required before configuring an interface.

Overview

In this example, you delete the t3-1/0/0 interface.

NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on the J-Web pages.

Configuration

IN THIS SECTION

- Procedure | 68

Procedure

Step-by-Step Procedure

To delete a T3 interface:

1. Specify the interface you want to delete.

```
[edit interfaces]
user@host# delete t3-1/0/0
```

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces` command.

Configuring 1-Port Clear Channel DS3/E3 GPIM

IN THIS SECTION

- Understanding the 1-Port Clear Channel DS3/E3 GPIM | 69
- Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode | 73

Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode | 76
Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode | 78

The 1-Port Clear Channel DS3/E3 GPIM is a channel interface that can support full-duplex DS3 (T3) or E3 line rates. The below topics shows the overview of the interface, examplse on how to configure the 1-Port Clear Channel DS3/E3 GPIM for DS3 port mode, E3 port mode and M23 mapping mode respectively.

Understanding the 1-Port Clear Channel DS3/E3 GPIM

IN THIS SECTION

Supported Features | 69
Interface Naming | 70
Physical Interface Settings | 70
Logical Interface Settings | 71

The 1-Port Clear Channel DS3/E3 Gigabit-Backplane *Physical Interface Module* (GPIM) for the device functions as a clear channel interface that can support full-duplex DS3 (T3) or E3 line rates of 44.796 or 34.368 Mbps, respectively. The DS3/E3 interface is a popular high-bandwidth WAN interface for large enterprise branch locations that enables high-quality voice, video, and data applications with reduced latency. The GPIM device does not support channelization, but it supports a subrate DS3/E3 configuration.

This topic includes the following sections:

Supported Features

The clear channel implementation provides such features as subrate and scrambling options used by major DSU vendors. The following key features are available depending on the interface and mode selections:

- Framed and unframed DS3 (default) and E3 port modes
- Support for frame relay, point-to-point, and HDLC serial encapsulation protocols
- Support for popular vendor algorithms for subrate and payload scrambling
- Support for generation and detection of loopback control codes (line-loopback activate and deactivate) and FEAC codes
- External and internal clocking support
- Support for DS3 and E3 network alarms
- Support for chassis clusters
- Support for anti-counterfeit check
- Loopback (local, remote, and payload) and BERT/PRBS/QRSS diagnostics support
- MTU size of 4474 bytes (default) and 9192 bytes (maximum)

Interface Naming

The following format represents the 1-Port Clear Channel DS3/E3 GPIM interface names:

```
type-fpc/pic/port
```

where:

- *type*—Media type (T3 or E3)
- *fpc*—Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- *pic*—Number of the PIC on which the physical interface is located
- *port*—Specific port on the PIC

Examples: `t3-1/0/0` and `e3-2/0/0`

Physical Interface Settings

The 1-Port Clear Channel DS3/E3 GPIM supports IP configurations. Using the CLI, you can configure the 1-Port Clear Channel DS3/E3 GPIM to operate in either DS3 or E3 mode. By default, at installation the physical interface, t3-x/y/z, is enabled on the GPIM port operating in DS3 mode with T3 framing.

You can reset the mode of the physical interface to E3 using the `edit chassis` command:

```
[edit]
user@host# set chassis fpc 1 pic 0 port 0 framing e3
```

Logical Interface Settings

The *logical interface* for the device is determined by setting the `t3-options` or `e3-options` of the `edit interfaces` command.

You can specify the MTU size for the GPIM interface. Junos OS supports an MTU value of 4474 bytes for the default value or up to 9192 bytes for maximum jumbo GPIM implementations.

Table 11 on page 71 identifies network interface specifications for DS3 or E3 modes.

Table 11: 1-Port Clear Channel DS3/E3 GPIM Interface Options

Description	DS3 Mode	E3 Mode
Network Interface Specifications		
Line encoding	B3ZS	HDB3
Framing	• C-bit parity (default) • M23	G.751 (default)
Subrate and scrambling	Vendor algorithms supported: • Adtran • Digital Link • Kentrox • Larscom • Verilink	Vendor algorithms supported: • Digital Link • Kentrox

Table 11: 1-Port Clear Channel DS3/E3 GPIM Interface Options *(Continued)*

Description	DS3 Mode	E3 Mode
Network alarms	Supported in accordance with the ANSI specification: • Loss of signal (LOS) • Out of frame (OOF) • Loss of frame (LOF) • Alarm identification Signal (AIS) • Remote defect identification (RDI)	Supported in accordance with the ITU-T specification: • Loss of signal (LOS) • Out of frame (OOF) • Alarm identification signal (AIS) • Remote defect identification (RDI) • Phase- locked loop (PLL)
Error counters	Incremented during a periodic 1-second polling routine: • Line code violations (LCV) • P-bit code violations (PCV) • C-bit code violations (CCV) • Line errored seconds (LES) • P-bit errored seconds (PES) • C-bit errored seconds (CES) • Severely errored framing seconds (SEFS) • P-bit severely errored seconds (PSES) • C-bit severely errored seconds (CSES) • Unavailable seconds (UAS)	Incremented during a periodic 1-second polling routine: • Frame alignment error (FAE) • Bipolar coding violations (BCV) • Excessive zeros (EXZ) • Line code violations (LCV) • Line errored seconds (LES) • Severely errored framing seconds (SEFS) • Unavailable seconds (UAS)
HDLC Features		

Table 11: 1-Port Clear Channel DS3/E3 GPIM Interface Options *(Continued)*

Description	DS3 Mode	E3 Mode
MTU	Default (4474 bytes) or maximum jumbo (up to 9192 bytes)	Default (4474 bytes) or maximum jumbo (up to 9192 bytes)
Shared flag	Supported	Supported
Idle flag/fill (0x7e or all ones)	Supported	Supported
Counters	Runts, giants	Runts, giants

SEE ALSO

Interface Naming Conventions | 9

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode

IN THIS SECTION

- Requirements | 73
- Overview | 74
- Configuration | 74

This example configures the GPIM in the DS3 (T3) operation mode.

Requirements

Before you begin:

- Install the device as specified in the *SRX Series Services Physical Interface Modules Hardware Guide*.

Overview

IN THIS SECTION

- Topology | 74

This example configures the basic T3 interface and modifies the framing to C-bit parity mode.

Topology

Configuration

IN THIS SECTION

- Procedure | 74

Procedure

Step-by-Step Procedure

To configure the GPIM:

1. Verify the installation, location, and status of the GPIM. In this example, the GPIM is installed in slot 8/PIC 0 and is currently online.

```
user@host> show chassis fpc pic-status
```

```
Slot 0 Online FPC
   PIC 0 Online 4x GE Base PIC
Slot 2 Offline FPC
Slot 5 Offline FPC
Slot 6 Online FPC
   PIC 0 Online 4x CT1E1 gPIM
Slot 7 Offline FPC
```

```
Slot 8 Online FPC
   PIC 0 Online 1x CLR CH T3/E3
```

2. Set the IP address for the logical interface.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet address interface 192.107.1.230/24
```

3. Set the MTU value to 9018.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet mtu 9018
```

4. Set the framing mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options cbit-parity
```

5. Enable the unframed DS3 mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options unframed
```

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

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

7. To verify the configuration for your device, enter the following operational command:

```
user@host> show interfaces t3-8/0/0 extensive
```

Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode

IN THIS SECTION

- Requirements | 76
- Overview | 76
- Configuration | 76

This example modifies the default configuration for an E3 environment.

Requirements

Before you begin:

- Install the device as specified in the *SRX Series Services Physical Interface Modules Hardware Guide*.

Overview

IN THIS SECTION

- Topology | 76

This example configures the basic E3 interface.

Topology

Configuration

IN THIS SECTION

- Procedure | 77

Procedure

Step-by-Step Procedure

To configure the GPIM in E3 framing:

1. Verify the installation, location, and status of the GPIM. In this example, the GPIM is installed in slot 8/PIC 0 and is currently online.

```
user@host> show chassis fpc pic-status
```

```
Slot 0 Online FPC
   PIC 0 Online 4x GE Base PIC
Slot 2 Offline FPC
Slot 5 Offline FPC
Slot 6 Online FPC
   PIC 0 Online 4x CT1E1 gPIM
Slot 7 Offline FPC
Slot 8 Online FPC
   PIC 0 Online 1x CLR CH T3/E3
```

2. Change to E3 port mode.

```
[edit]
user@host# set chassis fpc 8 pic 0 port 0 framing e3
```

3. Reset the MTU value to 3474.

```
[edit]
user@host# set interfaces e3-8/0/0 unit 0 family inet mtu 3474
```

4. Enable the unframed mode.

```
[edit]
user@host# set interfaces e3-8/0/0 e3-options unframed
```

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

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

6. To verify the configuration for your device, enter the following operational command:

```
user@host> show interfaces e3-8/0/0 extensive
```

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode

IN THIS SECTION

- Requirements | 78
- Overview | 78
- Configuration | 79

The following example configures the GPIM in DS3 with M23 mapping mode. Note that M23 mapping does not provide C-bit parity.

Requirements

Before you begin:

- Install the device as specified in the *SRX Series Services Physical Interface Modules Hardware Guide*.

Overview

IN THIS SECTION

- Topology | 79

This example configures the basic T3 interface and modifies the framing to M23 mode without C-bit parity.

Topology

Configuration

IN THIS SECTION

- Procedure | 79

Procedure

Step-by-Step Procedure

To configure the GPIM:

1. Verify the installation, location, and status of the GPIM. In this example, the GPIM is installed in slot 8/PIC 0 and is currently online.

```
user@host> show chassis fpc pic-status
```

```
Slot 0 Online FPC
   PIC 0 Online 4x GE Base PIC
Slot 2 Offline FPC
Slot 5 Offline FPC
Slot 6 Online FPC
   PIC 0 Online 4x CT1E1 gPIM
Slot 7 Offline FPC
Slot 8 Online FPC
   PIC 0 Online 1x CLR CH T3/E3
```

2. Set the IP address for the logical interface.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet address interface 192.107.1.230/24
```

3. Set the MTU value to 9018.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet mtu 9018
```

4. Set the framing mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options m23
```

5. Disable C-bit parity for M23 mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options no-cbit-parity
```

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

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

7. To verify the configuration for your device, enter the following operational command:

```
user@host> show interfaces t3-8/0/0 extensive
```

CHAPTER 3

Configuring ADSL and SHDSL Interfaces

ADSL and SHDSL Interfaces | 82

VDSL2 Interfaces | 125

ADSL and SHDSL Interfaces

SUMMARY

Learn about ADSL and SHDSL interface details and how to configure the interfaces on security devices.

IN THIS SECTION

- ADSL and SHDSL Interface Overview | **82**
- Example: Configure ADSL and SHDSL Network Interfaces | **86**
- Example: Configure G.SHDSL Interface | **106**

ADSL and SHDSL Interface Overview

Asymmetric digital subscriber line (ADSL) technology is part of the *x*DSL family of modem technologies that uses existing twisted-pair telephone lines to transport high-bandwidth data. The Symmetric high-speed DSL (SHDSL) interfaces support an SHDSL multirate technology which helps in data transfer between a single CPE subscriber and a central office (CO). The G.SHDSL Mini-*Physical Interface Module* (Mini-PIM) provides the physical connection to DSL network media types. Table 12 on page 82 specifies the key details of the ADSL, SHDSL interfaces, and G.SHDSL Mini-PIM.

Table 12: ADSL and SHDSL Interface Details

Interface Details	Description
Interface name	ADSL, SHDSL
Supported on	For information about platforms support, see hardware compatibility tool (HCT).
Interface type	• `at-` represents ADSL2, SHDSL interface and G.SHDSL Mini-PIM when you configure `at-` to function as VDSL2.

Table 12: ADSL and SHDSL Interface Details *(Continued)*

Interface Details	Description
ADSL/ADSL2/ ADSL2+ use cases	• Connects the loop between service provider networks and customer sites. ADSL Mini-PIM facilitates a maximum of 10 virtual circuits on supported security devices and can use PPP over Ethernet over ATM (PPPoEoA) and PPP over ATM (PPPoA) to connect through ADSL lines only. • Modems work as a dual-purpose ADSL circuit and can accommodate lower-frequency voice traffic and higher-frequency data traffic. • Improves the data rate and reach performance, diagnostics, standby mode, and interoperability of ADSL modems. It doubles the possible downstream data bandwidth, enabling rates of 20 Mbps on telephone lines shorter than 5000 feet (1.5 km). • Uses Seamless Rate Adaptation (SRA) to change the data rate of a connection during operation with no interruptions or bit errors and the ADSL2 transceiver detects changes in channel conditions with data transmission parameters.
SHDSL use cases	• Supports an SHDSL multirate technology for data transfer between a single customer premises equipment (CPE) subscriber and a central office (CO). ITU-T G.991.2 is the official standard for describing SHDSL, also known as G.SHDSL. • Delivers a bandwidth of up to 2.3 Mbps in symmetrical directions. Compatible with ADSL and therefore causes very little, if any, interference between cables and is deployed on a network similar to ADSL.
GSHDSL Mini-PIM use cases	Provides the physical connection to DSL network media types and extended ATM CoS functionality to cells across the network. By default, unspecified bit rate (UBR) is used because the bandwidth utilization is unlimited. You can define bandwidth utilization with sustained cell rate and burst tolerance.

For information on ADSL2 hardware specifications, see 1-Port ADSL2+ Mini-Physical Interface Module Network Interface Specifications.

Features Supported on the ADSL, ADSL2, and SHDSL Interface

Table 13 on page 84 describes the key features supported on ADSL2 and SHDSL interfaces.

Table 13: Key Features Supported on ADSL2 and SHDSL

Feature	Description
ADSL Features	
DSL	• Supports ATM-over-ADSL and ATM-over-SHDSL interfaces. Payload loopback functionality is not supported on ATM-over-SHDSL interfaces. • Uses PPP over Ethernet over ATM (PPPoEoA) and PPP over ATM (PPPoA) to connect through ADSL lines only for supported security devices with Mini-PIMs.
ATM CoS Support	Ability of a network to guarantee *class of service* depends on the way in which the source generates cells and on the availability of network resources. Based on the way in which the source generates cells and the availability of network resources, the set of traffic descriptors specified are: • Peak cell rate (PCR)—Top rate at which traffic can burst. • Sustained cell rate (SCR)—Normal traffic rate averaged over time. • Maximum burst size (MBS)—Maximum burst size that can be sent at the peak rate. • Cell delay variation tolerance (CDVT)—Allows you to delay the traffic for a particular time duration in microseconds to follow a rhythmic pattern.
Encapsulation	You can enable an existing Junos OS CLI to support MLPPP encapsulation and the family mlppp. To establish an ADSL link, you must first use an RJ-11 cable to connect the CPE to a DSLAM patch panel to form an ADSL link and then use OC3 or DS3 to connect the DSLAM to M Series or E Series devices to form an ATM backbone.
SHDSL Features	
Bandwidth	SHDSL is symmetrical and delivers a bandwidth of up to 2.3 Mbps in both directions. Compatible with ADSL and therefore causes very little, if any, interference between cables.

Table 13: Key Features Supported on ADSL2 and SHDSL *(Continued)*

Feature	Description
Packet Transfer Mode (PTM)	Supports PTM and packets (IP, PPP, Ethernet, MPLS, and so on) are transported over DSL links as an alternative to using Asynchronous Transfer Mode (ATM). PTM is based on the Ethernet in the First Mile (EFM) IEEE 802.3ah standard.
DSL	G.SHDSL Mini-*Physical Interface Module* (Mini-PIM) provides the physical connection to DSL network media types.
GSHDSL Virtual circuits (VC)	VC per Mini-PIM (10 maximum including OAM VC).
MTU size	Maximum MTU size of 9180 bytes.
GSHDSL PTM EFM	• Supports EFM PIC mode, PPPoE encapsulation, IPv6, *Chassis cluster* mode, and VLAN over EFM. • Maximum MTU size of 1514 bytes.

For more information on supported features and profiles on ADSL2 interfaces, see 1-Port ADSL2+ Mini-Physical Interface Module Key Features and for SHDSL and GSHDSL interfaces, see 1-Port G.SHDSL 8-Wire Mini-Physical Interface Module Overview.

ADSL transmission is asymmetric because the downstream bandwidth is typically greater than the upstream bandwidth. The typical bandwidths of ADSL, ADSL2, and ADSL2+ circuits are defined in Table 14 on page 85.

Table 14: Standard Bandwidths of DSL Operating Modes

Operating Modes	Upstream	Downstream
ADSL	800 Kbps—1Mbps	8 Mbps
ADSL2	1—1.5 Mbps	12—14 Mbps
ADSL2+	1—1.5 Mbps	24—25 Mbps

Table 14: Standard Bandwidths of DSL Operating Modes *(Continued)*

Operating Modes	Upstream	Downstream
ADSL2+ Annex M	2.5—3 Mbps	25 Mbps

Operating Modes and Line Rates of the G.SHDSL Mini-PIM

The G.SHDSL Mini-PIM supports 2-wire (4-port 2-wire) mode, 4-wire (2-port 4-wire) mode, 8-wire (1-port 8-wire) mode, and EFM mode. The default operating mode is 2x 4-wire for this G.SHDSL Mini-PIM. G.SHDSL is supported on all devices using the symmetrical WAN speeds shown in Table 15 on page 86.

Table 15: Symmetrical WAN Speeds

Modes	Symmetrical WAN Speed Using Annex A and B	Symmetrical WAN Speed Using Annex F and G
2-wire	2.3 Mbps	From 768 Kbps to 5.696 Mbps
4-wire	4.6 Mbps	From 1.536 Mbps to 11.392 Mbps
8-wire	9.2 Mbps	From 3.072 Mbps to 22.784 Mbps
EFM mode	2.3 Mbps	From 768 Kbps to 5.696 Mbps
NOTE: A maximum of 16 Mbps is supported on SRX210, SRX220, SRX240, and SRX550 devices.		

Example: Configure ADSL and SHDSL Network Interfaces

In this example you configure the ADSL and SHDSL interface on an SRX Series device which supports LFI through an MLPPP. To support MLPPP encapsulation and the family mlppp on the ADSL interface on an SRX Series device, you enable an existing Junos OS CLI. To establish an ADSL link between network devices, you must use some intermediate connections. First, use an RJ-11 cable to connect the CPE (for example, an SRX Series device) to a DSLAM patch panel to form an ADSL link. Then use OC3 or DS3 to connect the DSLAM to M Series or E Series devices to form an ATM backbone.

Table 16 on page 87 specifies the CLI quick configuration commands used for configuring ADSL and SHDSL interfaces.

Table 16: CLI Quick Configuration

Configuration Step	CLI Quick Configuration Commands
Configure the DHCP client on ADSL interface	set interfaces at-1/0/0 encapsulation ethernet-over-atm set interfaces at-1/0/0 atm-options vpi 2 set interfaces at-1/0/0 dsl-options operating-mode auto set interfaces at-1/0/0 unit 0 set interfaces at-1/0/0 unit 0 encapsulation ether-over-atm-llc set interfaces at-1/0/0 unit 0 vci 2.122 set interfaces at-1/0/0 unit 0 family inet set interfaces at-1/0/0 unit 0 family inet dhcp
Configure the IPv6 address on an ADSL interface	set interfaces at-1/0/0 encapsulation ethernet-over-atm set interfaces at-1/0/0 atm-options vpi 2 set interfaces at-1/0/0 unit 0 encapsulation ether-over-atm-llc set interfaces at-1/0/0 unit 0 vci 2.118 set interfaces at-1/0/0 unit 0 family inet6 address 13:13::1/64
Configure ATM-over-ADSL network interfaces	set interfaces at-2/0/0 atm-options vpi 25 oam-liveness up-count 200 down-count 200 set interfaces at-2/0/0 atm-options vpi 25 oam-period 100 set interfaces at-1/0/0 unit 0 shaping cbr set interfaces at-1/0/0 unit 0 shaping vbr peak 33000 set interfaces at-1/0/0 dsl-options operating-mode auto set interfaces at-1/0/0 encapsulation ethernet-over-atm set interfaces at-1/0/0 unit 3 encapsulation atm-nlpid oam-liveness up-count 200 down-count 200 set interfaces at-1/0/0 unit 3 oam-period 100 set interfaces at-1/0/0 unit 3 family inet set interfaces at-1/0/0 unit 3 vci 35
Configure CHAP on DSL interfaces	set access profile A-ppp-client client client1 chap-secret my-secret set interfaces at-3/0/0 unit 0 ppp-options chap access-profile A-ppp-client local-name A-at-3/0/0.0 passive

Table 16: CLI Quick Configuration *(Continued)*

Configuration Step	CLI Quick Configuration Commands
Configure ATM-over-SHDSL network interfaces	**set chassis fpc 6 pic 0 shdsl pic-mode 1-port-atm** **set interfaces at-2/0/0 atm-options vpi 25 oam-liveness up-count 200 down-count 200** **set interfaces at-2/0/0 atm-options vpi 25 oam-period 100** **set interfaces at-2/0/0 encapsulation ethernet-over-atm shdsl-options annex annex-a** **set interfaces at-2/0/0 encapsulation ethernet-over-atm shdsl-options line-rate auto** **set interfaces at-2/0/0 encapsulation ethernet-over-atm shdsl-options loopback local** **set interfaces at-2/0/0 encapsulation ethernet-over-atm shdsl-options snr-margin current 5 snext 5** **set interfaces at-2/0/0 unit 3 encapsulation atm-nlpid** **set interfaces at-2/0/0 unit 3 oam-liveness up-count 200 down-count 200** **set interfaces at-2/0/0 unit 3 oam-period 100** **set interfaces at-2/0/0 unit 3 oam-period 100** **set interfaces at-2/0/0 unit 3 vci 35**

Configure the DHCP client on ADSL interface

In this example, you configure the ATM interface as `at-1/0/0`. Then set the logical interface to `unit 0` and specify the family protocol type as `inet`. Finally, configure the DHCP client. To configure DHCP client on ADSL interfaces:

1. Set the encapsulation mode.

```
[edit]
user@host# set interfaces at-1/0/0 encapsulation ethernet-over-atm
```

2. Configure the ATM VPI option.

```
[edit]
user@host# set interfaces at-1/0/0 atm-options vpi 2
```

3. Set operating mode.

```
[edit]
user@host# set interfaces at-1/0/0 dsl-options operating-mode auto
```

4. Set the logical interface.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0
```

5. Set the encapsulation mode for logical interface.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 encapsulation ether-over-atm-llc
```

6. Set the ATM VCI option.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 vci 2.122
```

7. Specify the family protocol type.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet
```

8. Configure the DHCP client.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet dhcp
```

9. Set the DHCP client identifier as a ASCII or hexadecimal value (optional):

   Use hexadecimal if the client identifier is a MAC address—for example, 00:0a:12:00:12:12.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet dhcp client-identifier
00:0a:12:00:12:12
```

10. Set the DHCP lease time in seconds—for example, 86400 (24 hours). The range is 60 through 2147483647 seconds (optional).

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet dhcp lease-time 86400
```

11. Define the number of attempts allowed to retransmit a DHCP packet (optional)—for example, 6

    The range is 0 through 6. The default is 4 times.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet dhcp retransmission-attempt 6
```

12. Define the interval, in seconds, allowed between retransmission attempts (optional)—for example, 5.

    The range is 4 through 64. The default is 4 seconds.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet dhcp retransmission-interval 5
```

13. Set the IPv4 address of the preferred DHCP server (optional)—for example, 10.1.1.1.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet dhcp server-address 10.1.1.1
```

14. Set the vendor class ID for the DHCP client (optional)—for example, ether.

```
[edit]
user@host# set interfaces at-0/0/1 unit 0 family inet dhcp vendor-id ether
```

Use the `show interfaces at-1/0/0` command to see the output of the configuration.

Configure the IPv6 Address on an ADSL Interface

To configure the IPv6 address on an ADSL interface:

1. Configure the encapsulation type.

```
[edit]
user@host# set interfaces at-1/0/0 encapsulation ethernet-over-atm
```

2. Specify the annex type.

```
[edit]
user@host# set interfaces at-1/0/0 atm-options vpi 2
```

3. Configure the encapsulation for the logical unit.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 encapsulation ether-over-atm-llc
```

4. Configure the VCI value.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 vci 2.118
```

5. Configure family protocol type and assign an IPv6 address.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet6 address 13:13::1/64
```

Use the `show interfaces at-1/0/0` command to see the output of the configuration.

Configure ATM-over-ADSL Network Interfaces

This example shows how to use devices with ADSL Annex A or Annex B PIMs to send network traffic through a point-to-point connection to a DSLAM. Within the example, you set the DSL operating mode type to auto so that the ADSL interface will autonegotiate settings with the DSLAM.

The example shows how to create an ATM interface called at-2/0/0. The values for the interface's physical properties are kept relatively low—the ATM VPI is set to 25; both the OAM down count and up count are set to 200 cells; the OAM period is set to 100 seconds.

The example also shows how to set traffic shaping values on the ATM interface to support CoS. CBR is enabled in order to stabilize the cell transmission rate throughout the duration of the connection. Additionally, the VBR peak is set to 33,000 for data packet transfers.

Within the example, you set the encapsulation mode to ethernet-over-atm to support PPP over Ethernet IPv4 traffic. You also configure a logical interface (unit 3). The logical interface uses ATM NLPID encapsulation. As with the physical interface, the OAM down count and up count are set to 200 cells on the logical interface and the OAM period is set to 100 seconds. The family protocol is set to inet and the VCI is set to 35.

On SRX devices, the ATM interface takes more than 5 minutes to come up when CPE is configured in ANSI-DMT mode and CO is configured in automode. This occurs only with ALU 7300 DSLAM, due to limitation in current firmware version running on the ADSL Mini-PIM.

To configure ATM-over-ADSL network interfaces for the devices:

1. Create an ATM interface.

```
[edit]
user@host# edit interfaces at-2/0/0
```

2. Configure the physical properties for the ATM interface.

```
[edit interfaces at-2/0/0]
user@host# set atm-options vpi 25
user@host# set atm-options vpi 25 oam-liveness up-count 200 down-count 200
user@host# set atm-options vpi 25 oam-period 100
```

3. Specify the CBR value and VBR value for the Ethernet interface.

```
[edit]
user@host# edit interfaces at-1/0/0 unit 0
user@host# set shaping cbr
user@host# set shaping vbr peak 33000
```

4. Set the DSL operating mode type.

```
[edit interfaces at-1/0/0.0]
user@host# set dsl-options operating-mode auto
```

5. Configure the encapsulation type.

```
[edit interfaces at-1/0/0]
user@host# set encapsulation ethernet-over-atm
```

6. Configure the encapsulation for the logical unit.

```
[edit interfaces at-1/0/0 unit 3]
user@host# set encapsulation atm-nlpid
```

7. Configure the OAM liveness values for an ATM virtual circuit.

```
[edit interfaces at-1/0/0 unit 3]
user@host# set oam-liveness up-count 200 down-count 200
```

8. Specify the OAM period.

```
[edit interfaces at-1/0/0 unit 3]
user@host# set oam-period 100
```

9. Set the family protocol type.

```
[edit interfaces at-1/0/0 unit 3]
user@host# set family inet
```

10. Configure the VCI value.

```
[edit interfaces at-1/0/0 unit 3]
user@host# set vci 35
```

Use the `show` command to see the output of the configuration.

Configure MLPPP-over-ADSL Interfaces

In this example, you set the encapsulation as atm-mlppp-llc for the interface `at-5/0/0`. You then configure the family MLPPP bundle as `lsq-0/0/0.1`.

Figure 7 on page 94 shows a typical example of MLPPP-over-ADSL end-to-end connectivity.

Figure 7: MLPPP-over-ADSL Interface

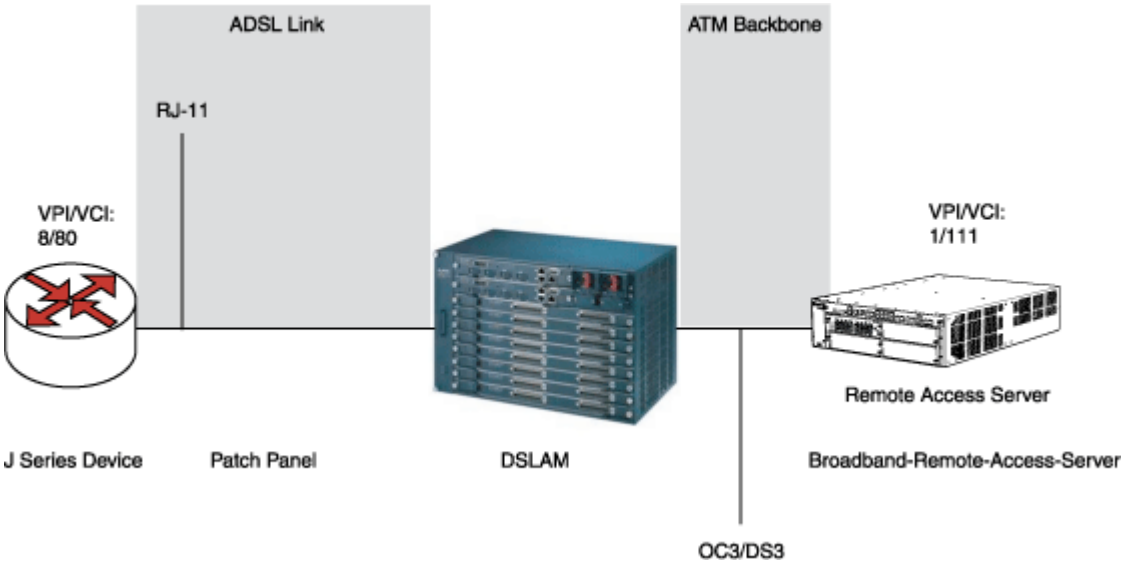

To configure MLPPP on an ADSL interface:

1. Configure an interface.

```
[edit]
user@host# edit interfaces at-5/0/0 unit 0
```

2. Set the MLPPP encapsulation.

```
[edit interfaces at-5/0/0 unit 0]
user@host# set encapsulation atm-mlppp-llc
```

3. Specify the family MLPPP.

```
[edit interfaces at-5/0/0 unit 0]
user@host# set family mlppp bundle lsq-0/0/0.1
```

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

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

Use the `show` command to see the output of the configuration.

Configure CHAP on DSL Interfaces

In this example, you specify the CHAP access profile and create an interface called `at-3/0/0`. You configure CHAP on either the ATM-over-ADSL or the ATM-over-SHDSL interface and specify a unique profile name called A-ppp-client containing a client list and access parameters. You then specify a unique hostname called `A-at-3/0/0.0` to be used in CHAP. Finally, you set the passive option to handle incoming CHAP packets. To configure CHAP on either the ATM-over-ADSL or the ATM-over-SHDSL interface:

1. Define a CHAP access profile.

```
[edit]
user@host# set access profile A-ppp-client client client1 chap-secret my-secret
```

2. Create an interface.

```
[edit]
user@host# edit interfaces at-3/0/0 unit 0
```

3. Configure CHAP and specify a unique profile name.

```
[edit interfaces at-3/0/0 unit 0]
user@host# set ppp-options chap access-profile A-ppp-client
```

4. Specify a unique hostname.

```
[edit interfaces at-3/0/0 unit 0]
user@host# set ppp-options chap local-name A-at-3/0/0.0
```

5. Set the option to handle incoming CHAP packets only.

```
[edit interfaces at-3/0/0 unit 0]
user@host# set ppp-options chap passive
```

Use the `show` command to see the output of the configuration.

Configure ATM-over-SHDSL Network Interfaces

In this example, you set the ATM-over-SHDSL mode on the G.SHDSL interface, if required. You create an interface called `at-2/0/0` and configure the physical properties for the interface. You configure the encapsulation type and annex type. You specify the SHDSL line rate for the ATM-over-SHDSL interface and the loopback address for testing the SHDSL connection integrity. Then you configure the SNR margin, set the logical interface, and configure the encapsulation for the ATM-over-SHDSL logical unit.

Additionally, you configure the OAM liveness values for an ATM virtual circuit and set the OAM period, Finally, you add the family protocol type inet and configure the VCI value. To configure ATM-over-SHDSL network interfaces for the device:

1. Set the ATM-over-SHDSL mode on the G.SHDSL interface.

```
[edit]
user@host# set chassis fpc 6 pic 0 shdsl pic-mode 1-port-atm
```

2. Create an interface.

```
[edit]
user@host# edit interfaces at-2/0/0
```

3. Configure the physical properties for the interface.

```
[edit interfaces at-2/0/0]
user@host# set atm-options vpi 25
user@host# set atm-options vpi 25 oam-liveness up-count 200 down-count 200
user@host# set atm-options vpi 25 oam-period 100
```

4. Configure the encapsulation type.

```
[edit interfaces at-2/0/0]
user@host# set encapsulation ethernet-over-atm
```

5. Set the annex type.

```
[edit]
user@host# edit interfaces at-2/0/0 shdsl-options
user@host# set annex annex-a
```

6. Configure the SHDSL line rate.

```
[edit interfaces at-2/0/0 shdsl-options]
user@host# set line-rate auto
```

7. Configure the loopback option for testing the SHDSL connection integrity.

```
[edit interfaces at-2/0/0 shdsl-options]
user@host# set loopback local
```

8. Configure the signal-to-noise ration margin.

```
[edit interfaces at-2/0/0 shdsl-options]
user@host# set snr-margin current 5
user@host# set snr-margin snext5
```

9. Configure the logical interface.

```
[edit]
user@host# edit interfaces at-2/0/0 unit 3
```

10. Configure the encapsulation for the logical unit.

```
[edit interfaces at-2/0/0 unit 3]
user@host# set encapsulation atm-nlpid
```

11. Configure the OAM liveness values for an ATM virtual circuit

```
[edit interfaces at-2/0/0 unit 3]
user@host# set oam-liveness up-count 200 down-count 200
```

12. Configure the OAM period.

```
[edit interfaces at-2/0/0 unit 3]
user@host# set oam-period 100
```

13. Add the Family protocol type.

```
[edit interfaces at-2/0/0 unit 3]
user@host# set family inet
```

14. Configure the VCI value.

```
[edit interfaces at-2/0/0 unit 3]
user@host# set vci 35
```

Use the `show` command to see the output of the configuration.

Verification

Display information about the parameters configured on the ADSL and SHDSL interfaces.

- To verify that the DHCP options are configured use the `run show system services dhcp client` command:

```
user@host# run show system services dhcp client

Logical Interface name         at-1/0/0.0
        Hardware address        00:1f:12:e4:71:38
        Client status           bound
```

```
        Address obtained        10.40.1.2
        Update server           disabled
        Lease obtained at       2011-05-03 04:58:10 PDT
        Lease expires at        2011-05-04 04:58:10 PDT

DHCP options:
    Name: server-identifier, Value: 10.40.1.1
    Code: 1, Type: ip-address, Value: 255.255.255.0
    Name: name-server, Value: [ 192.168.5.68, 192.168.60.131, 172.17.28.100, 172.17.28.101 ]
    Name: domain-name, Value: englab.juniper.net
```

To verify the interface status and check traffic statistics use the `show interface terse` command and test end-to-end data path connectivity by sending the ping packets to the remote end IP address:

```
user@host# run show interfaces at-1/0/0 terse

Interface               Admin Link Proto    Local                 Remote
at-1/0/0                up    up
at-1/0/0.0              up    up   inet     10.40.1.2/24
at-1/0/0.32767          up    up

user@host# run ping 10.40.1.1 count 100 rapid

PING 10.40.1.1 (10.40.1.1): 56 data bytes
!!!
```

```
            12:12::2         00:19:e2:4b:61:83      stale       1188 yes     no
at-3/0/0.0
```

The IPv6 Address field displays the configured IPv6 address on the interface.

- To verify the ADSL interface properties, use the show interfaces at-1/0/0 extensive command:

```
user@host> show interfaces at-1/0/0 extensive
 Physical interface: at-1/0/0, Enabled, Physical link is Up
   Interface index: 141, SNMP ifIndex: 49, Generation: 142
   Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, ADSL mode,
   Speed: ADSL, Loopback: None
   Device flags   : Present Running
   Link flags     : None
   CoS queues     : 8 supported, 8 maximum usable queues
   Hold-times     : Up 0 ms, Down 0 ms
   Current address: 00:05:85:c3:17:f4
   Last flapped   : 2008-06-26 23:11:09 PDT (01:41:30 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, Invalid VCs: 0, Framing errors: 0,Policed discards: 0, L3
incompletes: 0, L2 channelerrors: 0, L2 mismatch timeouts: 0,
     Resource errors: 0
   Output errors:
     Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0, Resource
errors: 0
   ADSL  alarms   : None
   ADSL  defects  : None
   ADSL media:            Seconds        Count State
    LOF                       1            1  OK
    LOS                       1            1  OK
    LOM                       0            0  OK
    LOP                       0            0  OK
    LOCDI                     0            0  OK
    LOCDNI                    0            0  OK
   ADSL status:
     Modem status  : Showtime (Adsl2plus)
```

```
     DSL mode      :     Auto    Annex A
     Last fail code: None
     Subfunction   : 0x00
     Seconds in showtime : 6093
   ADSL Chipset Information:              ATU-R               ATU-C
     Vendor Country :                      0x0f                0xb5
     Vendor ID      :                      STMI               IFTN
     Vendor Specific:                      0x0000             0x70de
   ADSL Statistics:                       ATU-R               ATU-C
     Attenuation (dB)         :            0.0                 0.0
     Capacity used(%)         :            100                 92
     Noise margin(dB)         :            7.5                 9.0
     Output power (dBm)       :            10.0                12.5
                                  Interleave       Fast  Interleave         Fast
     Bit rate (kbps)   :                  0      24465           0         1016
     CRC               :                  0          0           0         0
     FEC               :                  0          0           0         0
     HEC               :                  0          0           0         0
     Received cells    :                  0         49
     Transmitted cells :                  0          0
   ATM status:
     HCS state:     Hunt
     LOC      :       OK
  ATM Statistics:
    Uncorrectable HCS errors: 0, Correctable HCS errors: 0,Tx cell FIFO overruns: 0,Rx cell
FIFO overruns: 0,Rx cell FIFO underruns: 0,
    Input cell count: 49, Output cell count: 0,Output idle cell count: 0,Output VC queue
drops: 0Input no buffers: 0, Input length errors: 0,
    Input timeouts: 0, Input invalid VCs: 0, Input bad CRCs: 0, Input OAM cell no buffers: 0

  Packet Forwarding Engine configuration:
     Destination slot: 1
     Direction : Output
     CoS transmit queue               Bandwidth              Buffer  Priority   Limit
                              %            bps     %           usec
     0 best-effort            95        7600000    95              0       low
none
     3 network-control         5         400000     5              0       low    none



     But for ADSL MiniPim TI chipset does not send  ADSL Chipset
     Information. Also Adsl minipim does not send any alarms. So we can't
```

```
     show alarm stats for minipim. So following information will not be
     displayed in Minipim case.

     ADSL  alarms   : None
     ADSL  defects  : None
     ADSL media:            Seconds        Count State
     LOF                       1            1  OK
     LOS                       1            1  OK
     LOM                       0            0  OK
     LOP                       0            0  OK
     LOCDI                     0            0  OK
     LOCDNI                    0            0  OK

    ADSL Chipset Information:              ATU-R      ATU-C
      Vendor Country :                      0x0f       0xb5
      Vendor ID      :                      STMI       IFTN
      Vendor Specific:                      0x0000     0x70de
```

The output shows a summary of interface information.

To verify the PPPoA configuration for an ATM-over-ADSL interface is correct, use the the `show interfaces at-1/0/0` and the `show access` commands.

- To verify the configuration for an MLPPP-over-ADSL Interface is correct, use the `show interfaces at-5/0/0` command.
- To verify that the ADSL interface properties are enabled, use the `show interfaces at-3/0/0 extensive` command.

```
user@host> show interfaces at-3/0/0 extensive
 Physical interface: at-3/0/0, Enabled, Physical link is Up
   Interface index: 141, SNMP ifIndex: 49, Generation: 142
   Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, ADSL mode,
   Speed: ADSL, Loopback: None
   Device flags   : Present Running
   Link flags     : None
   CoS queues     : 8 supported, 8 maximum usable queues
   Hold-times     : Up 0 ms, Down 0 ms
   Current address: 00:05:85:c3:17:f4
   Last flapped   : 2008-06-26 23:11:09 PDT (01:41:30 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, Invalid VCs: 0, Framing errors: 0,Policed discards: 0, L3
incompletes: 0, L2 channelerrors: 0, L2 mismatch timeouts: 0,
     Resource errors: 0
   Output errors:
     Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0, Resource
errors: 0
   ADSL  alarms   : None
   ADSL  defects  : None
   ADSL media:            Seconds        Count State
    LOF                       1            1  OK
    LOS                       1            1  OK
    LOM                       0            0  OK
    LOP                       0            0  OK
    LOCDI                     0            0  OK
    LOCDNI                    0            0  OK
   ADSL status:
     Modem status  : Showtime (Adsl2plus)
     DSL mode      :     Auto    Annex A
     Last fail code: None
     Subfunction   : 0x00
     Seconds in showtime : 6093
   ADSL Chipset Information:              ATU-R               ATU-C
     Vendor Country :                      0x0f                0xb5
     Vendor ID      :                      STMI               IFTN
     Vendor Specific:                      0x0000             0x70de
   ADSL Statistics:                       ATU-R               ATU-C
     Attenuation (dB)         :            0.0                 0.0
     Capacity used(%)         :            100                 92
     Noise margin(dB)         :            7.5                 9.0
     Output power (dBm)       :            10.0                12.5
                                  Interleave       Fast  Interleave         Fast
     Bit rate (kbps)   :                  0      24465           0         1016
     CRC               :                  0          0           0         0
     FEC               :                  0          0           0         0
     HEC               :                  0          0           0         0
     Received cells    :                  0         49
     Transmitted cells :                  0          0
   ATM status:
     HCS state:     Hunt
```

```
     LOC      :       OK
  ATM Statistics:
    Uncorrectable HCS errors: 0, Correctable HCS errors: 0,Tx cell FIFO overruns: 0,Rx cell
FIFO overruns: 0,Rx cell FIFO underruns: 0,
    Input cell count: 49, Output cell count: 0,Output idle cell count: 0,Output VC queue
drops: 0Input no buffers: 0, Input length errors: 0,
    Input timeouts: 0, Input invalid VCs: 0, Input bad CRCs: 0, Input OAM cell no buffers: 0

  Packet Forwarding Engine configuration:
     Destination slot: 1
     Direction : Output
     CoS transmit queue               Bandwidth              Buffer  Priority   Limit
                               %            bps     %           usec
     0 best-effort            95        7600000    95              0       low
none
     3 network-control         5         400000     5              0       low    none



     But for ADSL MiniPim TI chipset does not send  ADSL Chipset
     Information. Also Adsl minipim does not send any alarms. So we can't
     show alarm stats for minipim. So following information will not be
     displayed in Minipim case.

     ADSL  alarms   : None
     ADSL  defects  : None
     ADSL media:            Seconds        Count State
     LOF                       1            1  OK
     LOS                       1            1  OK
     LOM                       0            0  OK
     LOP                       0            0  OK
     LOCDI                     0            0  OK
     LOCDNI                    0            0  OK

    ADSL Chipset Information:              ATU-R      ATU-C
      Vendor Country :                      0x0f       0xb5
      Vendor ID      :                      STMI       IFTN
      Vendor Specific:                      0x0000     0x70de
```

To verify the PPPoA configuration for an ATM-over-ADSL interface is correct, use the `show interfaces at-3/0/0` and the `show access` commands.

To verify that an ATM-over-SHDSL configuration is correct, use the `show interfaces at-3/0/0 extensive` command.

```
user@host> show interfaces at-3/0/0 extensive
Physical interface: at-3/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 23, Generation: 48
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, ADSL mode, Speed: ADSL,
  Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:c7:44:3c
  Last flapped   : 2005-05-16 05:54:41 PDT (00:41:42 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 4520                    0 bps
   Output bytes  :                39250                    0 bps
   Input  packets:                   71                    0 pps
   Output packets:                 1309                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 1, L2 mismatch timeouts: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    4                    4                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                2340                 2340                    0
  SHDSL  alarms   : None
  SHDSL  defects  : None
  SHDSL media:            Seconds        Count  State
    LOSD                  239206            2  OK
    LOSW                  239208            1  OK
    ES                         3            1  OK
    SES                        0            0  OK
    UAS                        3            1  OK

  SHDSL status:
    Line termination     :STU-R
```

```
    Annex                    :Annex B
    Line Mode                :2-wire
    Modem Status            :Data
    Last fail code        :0
    Framer mode            :ATM
    Dying Gasp               :Enabled
    Chipset version        :1
    Firmware version        :R3.0
SHDSL Statistics:
 Loop Attenuation (dB)            :0.600
  Transmit power (dB)                :8.5
  Receiver gain (dB)               :21.420
  SNR sampling (dB)                   :39.3690
  Bit rate     (kbps)                   :2304
  Bit error rate                       :0
  CRC errors                           :0
  SEGA errors                           :1
  LOSW errors                           :0
  Received cells                       :1155429
  Transmitted cells                   :1891375
  HEC errors                           :0
  Cell drop                           :0
```

Example: Configure G.SHDSL Interface

This example shows how to configure the G.SHDSL interface on SRX Series devices.

To configure GSHDSL interface:

1. Specify the wire mode on the G.SHDSL interface. The default wire mode is 4-wire (2-port, 4-wire).
2. Specify the annex type. The default annex type is `auto`.
3. Specify the SHDSL line rate (speed of transmission of data on the SHDSL connection). The default line rate is `auto`.
4. Specify the encapsulation type. The `pt-` interface does not require encapsulation types.
5. Configure the encapsulation type.

Before you begin:

- Configure the network interfaces as necessary. See "Understanding Ethernet Interfaces" on page 160.
- Install the G.SHDSL Mini-PIM in the first slot of the SRX210 chassis.
- Connect the SRX210 device to a DSLAM (IP DSLAM and ATM DSLAM).

Figure 8 on page 107 shows the topology for the G.SHDSL Mini-PIM operating in 2X4-wire mode.

Figure 8: G.SHDSL Mini-PIM Operating in 2X4-Wire Mode

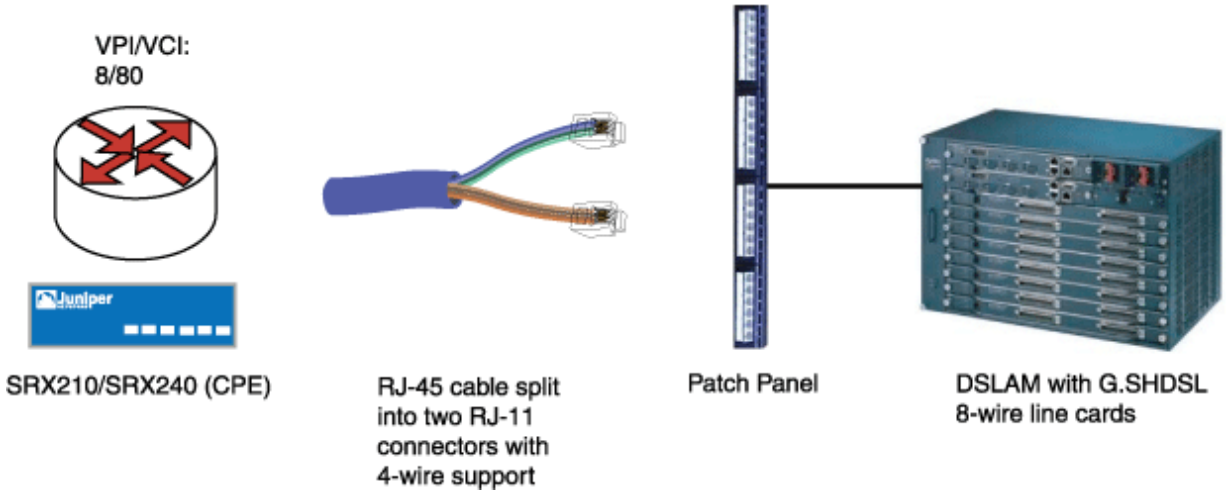

Figure 9 on page 107 shows the topology for the G.SHDSL Mini-PIM operating in 4X2-wire mode.

Figure 9: G.SHDSL Mini-PIM Operating in 4X2-Wire Mode

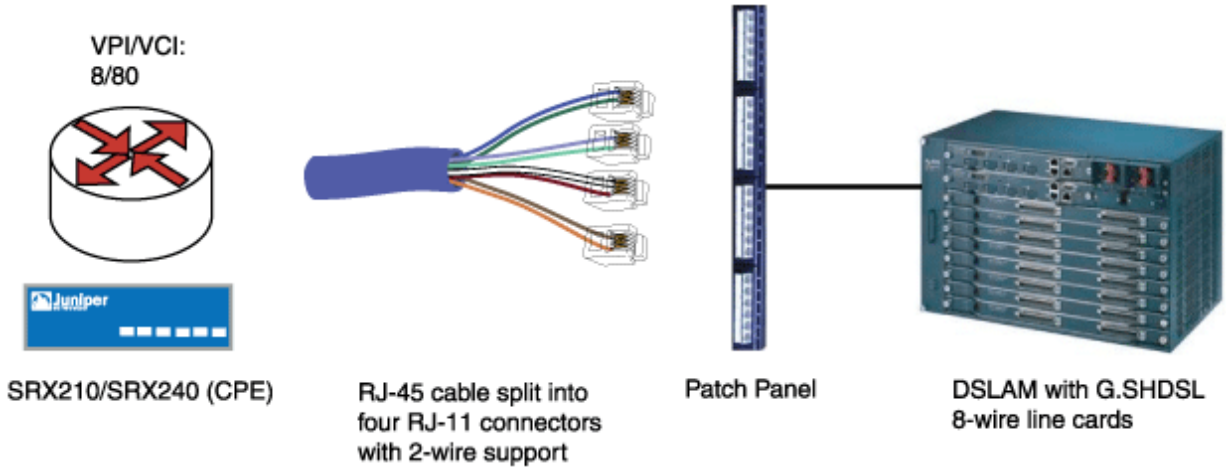

Figure 10 on page 108 shows the topology for the G.SHDSL Mini-PIM operating in 1X8-wire mode.

Figure 10: G.SHDSL Mini-PIM Operating in 1X8-Wire Mode

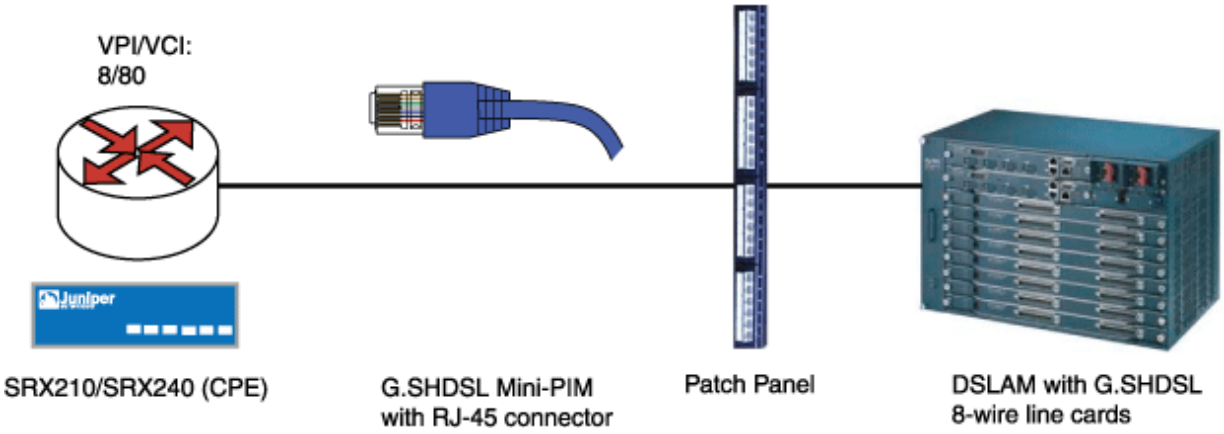

Determine the operating wire mode (2-wire, 4-wire, or 8-wire) and corresponding CLI code listed in Table 17 on page 108.

Table 17: Operating Wire Modes

Wire Mode Configuration	CLI Code
2x4-wire Configuration	`set chassis fpc 1 pic 0 shdsl pic-mode 2-port-atm` **NOTE**: The 2x4-wire configuration is the default configuration and behavior.
4x2-wire Configuration	`set chassis fpc 1 pic 0 shdsl pic-mode 4-port-atm`
1x8-wire Configuration	`set chassis fpc 1 pic 0 shdsl pic-mode 1-port-atm`

When you set the wire mode to 8-wire, one physical interface (IFD) is created. Similarly for 4-wire mode and 2-wire mode, two IFDs and four IFDs are created, respectively.

In this example:

1. First configure a basic G.SHDSL interface. Set the operation wire mode to 2-port-atm, the line rate to 4096, and the annex type to `annex-a`.
2. Configure the G.SHDSL interface when the device is connected to an IP DSLAM. Set the type of encapsulation to ethernet-over-atm and the ATM VPI option to 0. Set the type of encapsulation on

the G.SHDSL logical interface as `ether-over-atm-llc` and configure the ATM VCI option to 0.60. Also, set the interface address for the logical interface to 10.1.1.1/24.

3. Configure the G.SHDSL interface when the device is connected to an ATM DSLAM. Set the ATM VPI to 0 and set the type of encapsulation to `ppp-over-ether-over-atm-llc`. Specify a PPPoE interface with the PAP access profile, `local-name`, and `local-password`. Configure the passive option to handle incoming PAP packets and set the logical interface as the underlying interface for the PPPoE session to `at-1/0/0.0`. Set the number of seconds to 120 to wait before reconnecting after a PPPoE session is terminated. Specify the logical interface as the client for the PPPoE interface and obtain an IP address by negotiation with the remote end.

4. Configure PPPoA over ATM for the G.SHDSL Interface. set the type of encapsulation to `atm-pvc` and the ATM VPI to 0. Set the type of encapsulation for PPP over ATM adaptation layer 5 (AAL5) logical link control (LLC) on the logical interface and set the ATM VCI to 122. Configure the PPPoA interface with the CHAP access profile as juniper and set the `local-name` for the CHAP interface to srx-210. Finally, you obtain an IP address by negotiation with the remote end.

Table 18 on page 109specifies the CLI quick configuration commands used for configuring GSHDSL interfaces.

Table 18: CLI Quick Configuration

Configuration Step	CLI Quick Configuration Commands
Configure a basic G.SHDSL interface	`set chassis fpc 1 pic 0 shdsl pic-mode 2-port-atm` `set interfaces at-1/0/0 shdsl-options line-rate 4096 annex annex-a`

Configure G.SHDSL interface when connected to an IP DSLAM	`set interfaces at-1/0/0 encapsulation ethernet-over-atm`
	`set interfaces at-1/0/0 atm-options vpi 0`
	`set interfaces at-1/0/0 unit 0 encapsulation ether-over-atm-llc vci 0.60`
	`set interfaces at-1/0/0 unit 0  family inet address 10.1.1.1/24`

Configure G.SHDSL Interface when connected to an ATM DSLAM	`set interfaces at-1/0/0 encapsulation atm-pvc atm-options vpi 0`
	`set interfaces at-1/0/0 unit 0 encapsulation atm-snap vci 0.65`
	`set interfaces at-1/0/0 unit 0 family inet address 10.2.1.1/24`

Table 18: CLI Quick Configuration *(Continued)*

Configuration Step	CLI Quick Configuration Commands

Configure PPPoE over ATM for the G.SHDSL interface	`set interfaces at-1/0/0 encapsulation ethernet-over-atm  atm-options vpi 0`
	`set interfaces at-1/0/0 unit 0 encapsulation ppp-over-ether-over-atm-llc vci 0.35`
	`set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name srx-210`
	`set interfaces pp0 unit 0 ppp-options pap local-password  "$9$0tLw1SeN-woJDSr-wY2GU69Cp1RSre"`
	`set interfaces pp0 unit 0 ppp-options pap passive`
	`set interfaces pp0 unit 0 pppoe-options underlying-interface at-1/0/0.0`
	`set interfaces pp0 unit 0 pppoe-options auto-reconnect 120 client`
	`set interfaces pp0 unit 0 family inet negotiate-address`

Configure PPPoA over ATM for the G.SHDSL interface	`set interfaces at-1/0/0 encapsulation atm-pvc atm-options vpi 0`
	`set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc vci 1.122`
	`set interfaces at-1/0/0 unit 0 ppp-options chap access-profile juniper local-name srx-210`
	`set interfaces at-1/0/0 unit 0 family inet negotiate-address`

Cofigure a basic G.SHDSL interface in EFM PIC mode	`set chassis fpc 1 pic 0 shdsl pic-mode efm`
	`set interfaces pt-1/0/0 shdsl-options annex annex-g`
	`set interfaces pt-1/0/0 shdsl-options line-rate 5696`
	`set interfaces pt-1/0/0 unit 0 family inet address 10.10.10.1/24`

Configure PPPoE and VLAN for the G.SHDSL EFM interface	`set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether`
	`set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name srx-210`
	`set interfaces pp0 unit 0 ppp-options pap local-password  "$9$0tLw1SeN-woJDSr-wY2GU69Cp1RSre"`
	`set interfaces pp0 unit 0 ppp-options pap passive`
	`set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0`
	`set interfaces pp0 unit 0 pppoe-options auto-reconnect 120 client`
	`set interfaces pp0 unit 0 family inet negotiate-address`

Table 18: CLI Quick Configuration *(Continued)*

Configuration Step	CLI Quick Configuration Commands

Configure a Basic G.SHDSL Interface in EFM PIC Mode	`set chassis fpc 1 pic 0 shdsl pic-mode efm`
	`set interfaces pt-1/0/0 shdsl-options annex annex-g`
	`set interfaces pt-1/0/0 shdsl-options line-rate 5696`
	`set interfaces pt-1/0/0 unit 0 family inet address 10.10.10.1/24`

Configure PPPoE and VLAN for the G.SHDSL EFM Interface	`set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether`
	`set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name srx-210`
	`set interfaces pp0 unit 0 ppp-options pap local-password "$9$0tLw1SeN-woJDSr-wY2GU69Cp1RSre"`
	`set interfaces pp0 unit 0 ppp-options pap passive`
	`set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0`
	`set interfaces pp0 unit 0 pppoe-options auto-reconnect 120 client`
	`set interfaces pp0 unit 0 family inet negotiate-address`

Configure the Basic G.SHDSL Interfaces

To view the CLI quick configuration commands, see Table 18 on page 109. To configure the basic G.SHDSL interface on SRX210 devices:

1. Select the operating wire mode.

```
[edit]
user@host# set chassis fpc 1 pic 0 shdsl pic-mode 2-port-atm
```

2. Create an interface and set options.

```
[edit]
user@host# edit interfaces at-1/0/0 shdsl-options
```

3. Configure the line rates.

```
[edit interfaces at-1/0/0 shdsl-options]
user@host# set line-rate 4096
```

4. Set the annex type.

```
[edit interfaces at-1/0/0 shdsl-options]
user@host# set annex annex-a
```

Use the `show interfaces at-1/0/0` command to see the output of the configuration.

Configure a G.SHDSL Interface When Connected to an IP DSLAM

To configure the G.SHDSL interface on an SRX210 device when the device is connected to an IP DSLAM: :

1. Create an interface.

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

2. Specify the type of encapsulation.

```
[edit interfaces at-1/0/0]
user@host# set encapsulation ethernet-over-atm
```

3. Configure the ATM VPI option.

```
[edit interfaces at-1/0/0]
user@host# set atm-options vpi 0
```

4. Specify the type of encapsulation for logical interface.

```
[edit interfaces at-1/0/0 ]
user@host# edit unit 0
user@host# set encapsulation ether-over-atm-llc
```

5. Configure the ATM VCI options for the logical interface.

```
[edit interfaces at-1/0/0 unit 0]
user@host# set vci 0.60
```

6. Configure the interface address.

```
[edit interfaces at-1/0/0 unit 0]
user@host# set family inet address 10.1.1.1/24
```

Use the `show interfaces at-1/0/0` command to see the output of the configuration.

Configure a G.SHDSL Interface When Connected to an ATM DSLAM

To configure the G.SHDSL interface on an SRX210 device when the device is connected to an ATM DSLAM: :

1. Create an interface.

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

2. Specify the type of encapsulation.

```
[edit interfaces at-1/0/0]
user@host# set encapsulation atm-pvc
```

3. Configure the ATM VPI option.

```
[edit interfaces at-1/0/0]
user@host# set atm-options vpi 0
```

4. Specify the type of encapsulation for the logical interface.

```
[edit interfaces at-1/0/0]
user@host# edit unit 0
user@host# set encapsulation atm-snap
```

5. Configure the ATM VCI option.

```
[edit interfaces at-1/0/0 unit 0]
user@host# set vci 0.65
```

6. Configure the interface address.

```
[edit interfaces at-1/0/0 unit 0]
user@host# set family inet address 10.2.1.1/24
```

Use the `show interfaces at-1/0/0` command to see the output of the configuration.

Configure PPPoE over ATM for the G.SHDSL Interface

To configure PPPoE over ATM on the G.SHDSL interface:

1. Create an interface.

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

2. Specify the type of encapsulation.

```
[edit interfaces at-1/0/0]
user@host# set encapsulation ethernet-over-atm
```

3. Configure the ATM VPI option.

```
[edit interfaces at-1/0/0]
user@host# set atm-options vpi 0
```

4. Specify the type of encapsulation on the logical interface.

```
[edit interfaces at-1/0/0]
user@host# edit unit 0
user@host# set encapsulation ppp-over-ether-over-atm-llc
```

5. Configure the ATM VCI option.

```
[edit interfaces at-1/0/0 unit 0]
user@host# set vci 0.35
```

6. Configure a PPPoE interface with the PAP access profile.

```
[edit]
user@host# edit interfaces pp0 unit 0 ppp-options pap
user@host# set access-profile pap_prof
```

7. Configure a local-name for the PAP interface.

```
[edit interfaces pp0 unit 0 ppp-options pap]
user@host# set local-name srx-210
```

8. Configure a local-password for the PAP interface.

```
[edit interfaces pp0 unit 0 ppp-options pap]
user@host# set local-password "$9$0tLw1SeN-woJDSr-wY2GU69Cp1RSre"
```

9. Set the passive option to handle incoming PAP packets.

```
[edit interfaces pp0 unit 0 ppp-options pap]
user@host# set passive
```

10. Specify the logical interface as the underlying interface for the PPPoE session.

```
[edit]
user@host# edit interfaces pp0 unit 0 pppoe-options
user@host# set underlying-interface at-1/0/0.0
```

11. Specify the number of seconds.

```
[edit interfaces pp0 unit 0 pppoe-options]
user@host# set auto-reconnect 120
```

12. Set the logical interface as the client for the PPPoE interface.

```
[edit interfaces pp0 unit 0 pppoe-options]
user@host# set client
```

13. Obtain an IP address by negotiation with the remote end.

```
[edit]
user@host# edit interfaces pp0 unit 0
user@host# set family inet negotiate-address
```

Use the `show interfaces at-1/0/0` and `show interfaces pp0` commands to see the output of the configuration.

Configure PPPoA over ATM for the G.SHDSL Interface

To configure PPPoA over ATM on the G.SHDSL interface:

1. Create an interface.

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

2. Specify the type of encapsulation.

```
[edit interfaces at-1/0/0]
user@host# set encapsulation atm-pvc
```

3. Configure the ATM VPI option.

```
[edit interfaces at-1/0/0]
user@host# set atm-options vpi 0
```

4. Specify the type of encapsulation on the G.SHDSL logical interface.

```
[edit]
user@host# edit interfaces at-1/0/0 unit 0
user@host# set encapsulation atm-ppp-llc
```

5. Configure the ATM VCI option.

```
[edit interfaces at-1/0/0 unit 0]
user@host# set vci 1.122
```

6. Configure a PPPoA interface with the CHAP access profile.

```
[edit]
user@host# edit interfaces at-1/0/0 unit 0 ppp-options chap
user@host# set access-profile juniper
```

7. Configure a local name for the CHAP interface.

```
[edit interfaces at-1/0/0 unit 0 ppp-options chap]
user@host# set local-name srx-210
```

8. Obtain an IP address by negotiation with the remote end.

```
[edit]
user@host# edit interfaces at-1/0/0 unit 0
user@host# set family inet negotiate-address
```

Use the `show interfaces at-1/0/0` command to see the output of the configuration.

Configure G.SHDSL Interface in EFM Mode

In this example:

1. You first configure a basic G.SHDSL interface by setting the operation wire mode to efm, the line rate to auto, and the annex type to annex-auto.
2. You then configure the G.SHDSL interface when the device is connected to an EFM IP DSLAM. You set the logical interface to 10.10.10.1/24.
3. Next you configure PPPoE for the G.SHDSL Interface. Configure the encapsulation as ppp-over-ether under `unit 0` of `pt-1/0/0` interface. You specify a PPPoE interface with the PAP access profile, local name, and local password. Then you configure the passive option to handle incoming PAP packets and set the logical interface as the underlying interface for the PPPoE session to `pt-1/0/0.0`. Also, you set the number of seconds to 120 to wait before reconnecting after a PPPoE session is terminated. (The range is 1 through 4,294,967,295 seconds.) Finally, you specify the logical interface as the client for the PPPoE interface and obtain an IP address by negotiation with the remote end.

Figure 11 on page 118 shows the topology for the G.SHDSL Mini-PIM operating in EFM mode.

Figure 11: G.SHDSL Mini-PIM Operating in EFM Mode

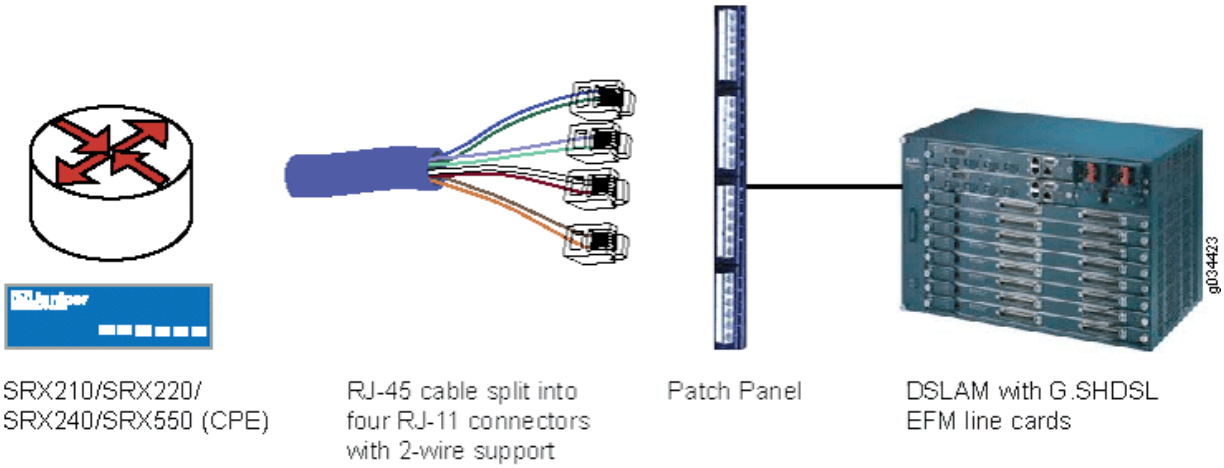

For operating wire mode EFM configuration, use the `set chassis fpc 1 pic 0 shdsl pic-mode efm` CLI code. When PIC mode is set to EFM, an interface called `pt-1/0/0` is created.

To view the CLI quick configuration commands, see Table 18 on page 109.

Configure a Basic G.SHDSL Interface in EFM PIC Mode

To configure a basic G.SHDSL interface:

1. Specify the PIC mode.

```
[edit]
user@host# set chassis fpc 1 pic 0 shdsl pic-mode efm
```

> **NOTE**: When configuring the G.SHDSL interface in chassis cluster mode, include the node ID. For example, to configure the G.SHDSL interface (operating in EFM PIC mode) in chassis cluster mode for fpc slot 1 on node 0, use the following command:
>
> **set chassis node 0 fpc 1 pic 0 shdsl pic-mode efm**

2. Configure the IP address.

```
[edit]
user@host# set interfaces pt-1/0/0 unit 0 family inet address 10.10.10.1/24
```

NOTE: By default, annex mode and line rate are set to auto. If you have to configure annex mode (annex-g) and line rate (5696 Kbps), follow Steps 3, 4, and 5.

3. Configure SHDSL options.

```
[edit]
user@host# set interfaces pt-1/0/0 shdsl-options
```

4. Specify the annex type.

```
[edit interfaces pt-1/0/0 shdsl-options]
user@host# set annex annex-g
```

5. Configure the line rate.

```
[edit interfaces pt-1/0/0 shdsl-options]
user@host# set line-rate 5696
```

Use the `show interfaces pt-1/0/0` and `show chassis fpc 1` commands to see the output of the configuration.

Configure a PPPoE and VLAN for the G.SHDSL EFM Interface

To configure PPPoE and VLAN for the G.SHDSL EFM Interface:

1. Create an interface.

```
[edit]
user@host# set interfaces pt-1/0/0
```

2. Specify the type of encapsulation.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0
user@host# set encapsulation ppp-over-ether
```

3. Configure a PPPoE interface with the PAP access profile.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap
user@host# set access-profile pap_prof
```

4. Configure a local name for the PAP interface.

```
[edit interfaces pp0 unit 0 ppp-options pap]
user@host# set local-name srx-210
```

5. Configure a local password for the PAP interface.

```
[edit interfaces pp0 unit 0 ppp-options pap]
user@host# set local-password "$9$0tLw1SeN-woJDSr-wY2GU69Cp1RSre"
```

6. Set the passive option to handle incoming PAP packets.

```
[edit interfaces pp0 unit 0 ppp-options pap]
user@host# set passive
```

7. Specify the logical interface as the underlying interface for the PPPoE session.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options
user@host# set underlying-interface pt-1/0/0.0
```

8. Specify the number of seconds.

```
[edit interfaces pp0 unit 0 pppoe-options]
user@host# set auto-reconnect 120
```

9. Set the logical interface as the client for the PPPoE interface.

```
[edit interfaces pp0 unit 0 pppoe-options]
user@host# set client
```

10. Obtain an IP address by negotiation with the remote end.

```
[edit interfaces]
user@host# set pp0 unit 0 family inet negotiate-address
```

11. Configure VLAN on EFM.

```
[edit interfaces]
user@host# set pt-1/0/0 vlan-tagging
```

12. Specify the VLAN ID.

```
[edit interfaces]
user@host# set pt-1/0/0 unit 0 vlan-id 99
```

Use the `show interfaces pt-1/0/0` and `show interfaces pp0` commands to see the output of the configuration.

Verification

Display information about the parameters configured on the GSHDSL interfaces.

- To display information about all the basic G.SHDSL interface properties, use the `show interfaces at-1/0/0 extensive` command.
- To display information about G.SHDSL interface properties:

```
user@host> show interfaces pt-1/0/0 extensive
```

EFM mode for interface pt-1/0/0:

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 575, Generation: 277
  Link-level type: Ethernet, MTU: 1514, Speed: SHDSL(8-Wire)
  Device flags   : Present Running
```

```
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 78:fe:3d:60:2f:99
  Last flapped   : 2012-10-11 00:03:13 PDT (00:28:57 ago)
  Statistics last cleared: 2012-10-11 00:32:05 PDT (00:00:05 ago)
  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, Invalid VCs: 0, Framing errors: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0, Resource
errors: 0
  EFM Group Statistics:
    Type                  : EFM bond
    Active Pairs          : 4
    Bit rate (in Kbps)    : 22784
  Line Pair  0  : Up
    Active alarms  : None
    Active defects : None
    SHDSL media:      Seconds   Count   State
      ES                    0
      SES                   0
      UAS                   0
    SHDSL status:
      Line  termination      : STU-R
      Annex                  : Annex G
      Line mode              : 2-wire
      Modem status           : Data
      Bit rate (kbps)        : 5696
      Last fail mode         : No failure (0x00)
      Framer mode            : EFM
      PAF Status             : Active
      Dying gasp             : Enabled
      Framer sync status     : In sync
    SHDSL statistics:
      Loop attenuation (dB)  : 0.0
      Transmit power (dBm)   : 14.0
      SNR sampling (dB)      : 14.0000
```

```
      CRC errors             : 2
      SEGA errors            : 0
      LOSW errors            : 0
  Line Pair  1  : Up
    Active alarms  : None
    Active defects : None
    SHDSL media:      Seconds   Count   State
      ES                    0
      SES                   0
      UAS                   0
    SHDSL status:
      Line  termination      : STU-R
      Annex                  : Annex G
      Line mode              : 2-wire
      Modem status           : Data
      Bit rate (kbps)        : 5696
      Last fail mode         : No failure (0x00)
      Framer mode            : EFM
      PAF Status             : Active
      Dying gasp             : Enabled
      Framer sync status     : In sync
    SHDSL statistics:
      Loop attenuation (dB)  : 0.0
      Transmit power (dBm)   : 14.0
      SNR sampling (dB)      : 19.0000
      CRC errors             : 0
      SEGA errors            : 0
      LOSW errors            : 0
  Line Pair  2  : Up
    Active alarms  : None
    Active defects : None
    SHDSL media:      Seconds   Count   State
      ES                    0
      SES                   0
      UAS                   0
    SHDSL status:
      Line  termination      : STU-R
      Annex                  : Annex G
      Line mode              : 2-wire
      Modem status           : Data
      Bit rate (kbps)        : 5696
      Last fail mode         : No failure (0x00)
      Framer mode            : EFM
```

```
      PAF Status             : Active
      Dying gasp             : Enabled
      Framer sync status     : In sync
    SHDSL statistics:
      Loop attenuation (dB)  : 0.0
      Transmit power (dBm)   : 14.0
      SNR sampling (dB)      : 14.0000
      CRC errors             : 0
      SEGA errors            : 0
      LOSW errors            : 0
  Line Pair  3  : Up
    Active alarms  : None
    Active defects : None
    SHDSL media:      Seconds   Count   State
      ES                    0
      SES                   0
      UAS                   0
    SHDSL status:
      Line  termination      : STU-R
      Annex                  : Annex G
      Line mode              : 2-wire
      Modem status           : Data
      Bit rate (kbps)        : 5696
      Last fail mode         : No failure (0x00)
      Framer mode            : EFM
      PAF Status             : Active
      Dying gasp             : Enabled
      Framer sync status     : In sync
    SHDSL statistics:
      Loop attenuation (dB)  : 1.0
      Transmit power (dBm)   : 14.0
      SNR sampling (dB)      : 18.0000
      CRC errors             : 0
      SEGA errors            : 0
      LOSW errors            : 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       21644800    95             0       low    none
    3 network-control          5        1139200     5             0       low    none
```

The output shows a summary of interface information.

RELATED DOCUMENTATION

- Understanding Point-to-Point Protocol over Ethernet | 308
- *Using the CLI Editor in Configuration Mode*
- Configuring the inet6 IPv6 Protocol Family | 36
- show ipv6 neighbors | 1047
- clear ipv6 neighbors | 781

VDSL2 Interfaces

SUMMARY

Learn about VDSL2 interface details and how to configure the interfaces on security devices.

IN THIS SECTION

- VDSL2 Interface Overview | 125
- Example: Configure VDSL2 Interface | 129

VDSL2 Interface Overview

IN THIS SECTION

- Features Supported on the VDSL2 Interface | 126
- VDSL2 Network Deployment Topology | 128

Very-high-bit-rate digital subscriber line (VDSL) technology is part of the xDSL family of modem technologies, which provide faster data transmission over a single flat untwisted or twisted pair of copper wires. Table 19 on page 126 specifies the key details of the VDSL2 interface.

Table 19: VDSL2 Interface Details

Interface Details	Description
Interface name	SRX-MP-1VDSL2-R
Supported on	For information about platforms support, see hardware compatibility tool (HCT).
Interface type	• `pt-` represents VDSL2 interface when you configure pt- to function as VDSL2. • Interface `pt-1/0/0` comes up by default.
Use cases	• Connects you and the service provider networks over a single connection to provide high bandwidth applications (triple-play services) like high-speed Internet access, Telephone services (VoIP (Voice over IP protocol), High-Definition TV (HDTV)), and Interactive gaming services. • VDSL2 carries the data and multimedia on the copper wire without interrupting the line's ability to carry voice signals. VDSL2 provides an ADSL interface in an ATM DSLAM topology and a VDSL2 interface in an IP or VDSL DSLM topology.

For information on VDSL2 hardware specifications, see 1-Port VDSL2 Annex A Mini-Physical Interface Module (SRX-MP-1VDSL2-R).

Features Supported on the VDSL2 Interface

Table 20 on page 126 describes the key features supported on VDSL2 interface.

Table 20: Key Features Supported on VDSL2

Feature	Description
Packet Transfer Mode (PTM)	• Uses the named interface `pt-1/0/0` and transports packets (IP, PPP, Ethernet, MPLS, and so on) over DSL links as an alternative to using Asynchronous Transfer Mode (ATM). • Based on the Ethernet in the First Mile (EFM) IEEE802.3ah standard.

Table 20: Key Features Supported on VDSL2 *(Continued)*

Feature	Description
Discrete multitone (DMT) modulation	• Separates a digital subscriber line signal to a usable frequency range of 256 frequency bands (or channels) with 4.3125 KHz each. • Uses the Fast Fourier Transform (FFT) algorithm for demodulation or modulation for increased speed.
Backward compatibility	• Backward compatible with most ADSL interface standards. • In ADSL fallback mode, VDSL2 operates on the ATM encapsulation interface in the first mile and uses the interface `at-1/0/0`. • Takes about 60 seconds to switch from VDSL2 to ADSL or from ADSL to VDSL2 operating modes.
Vectoring	• Employs coordination of line signals to reduce crosstalk levels to provide improved performance. • The ITU-T G.993.5 standard also known as G.vector, describes vectoring for VDSL2.
IPv6 Support	• Supports IPv6 on the DSL encapsulations like ATM physical interface encapsulations, atm-pvc, ethernet-over-atm, ethernet-over-atm, and ATM logical interface encapsulations except for atm-vc-mux and ppp-over-ether-over-atm-llc. • To configure IPv6 addresses on DSL interfaces in ATM or PTM mode, include the family protocol type as `inet6`.

For more information on supported features and profiles on VDSL2 interfaces, see1-Port VDSL2 Annex A Mini-Physical Interface Module (SRX-MP-1VDSL2-R).

VDSL2 Network Deployment Topology

The VDSL2 interface uses either Gigabit Ethernet or fiber as second mile to connect to the Broadband Remote Access Server (B-RAS). Figure 12 on page 128 shows a typical VDSL2 network topology.

Figure 12: Typical VDSL2 End-to-End Connectivity and Topology Diagram

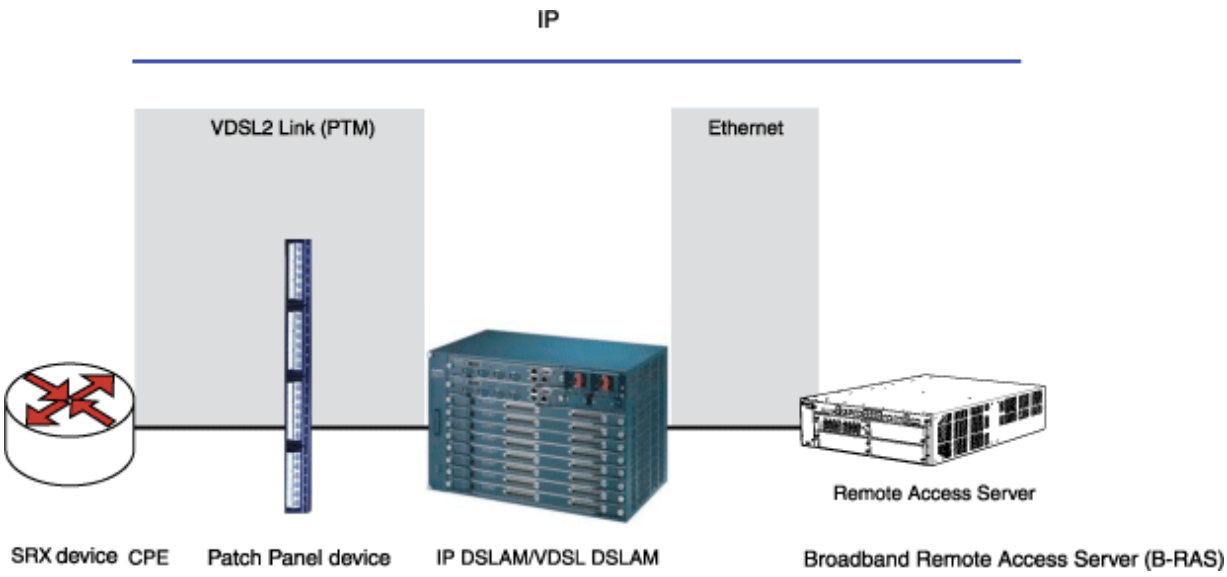

The ADSL interface uses either Gigabit Ethernet or OC3/DS3 ATM as the second mile to connect to the B-RAS. Figure 13 on page 128 shows a backward-compatible ADSL topology using ATM DSLAM.

Figure 13: Backward-Compatible ADSL Topology (ATM DSLAM)

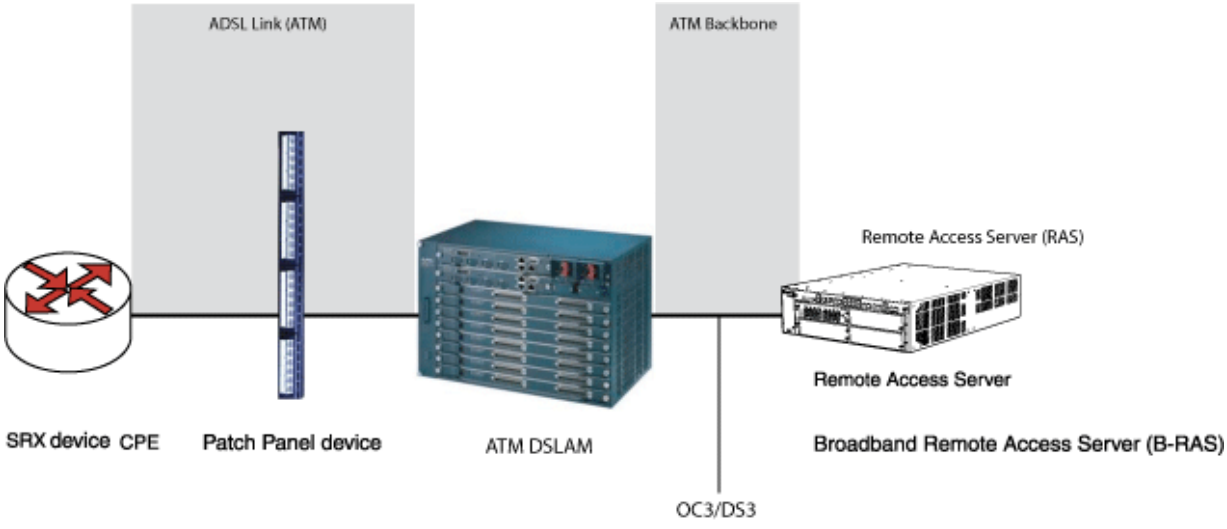

Example: Configure VDSL2 Interface

IN THIS SECTION

- Configure the VDSL2 Interface and Enable VLAN Tagging | 133
- Configure VDSL2 Interface with VDSL2 Mini-PIMs | 135
- Verification | 142

In this example you configure the VDSL2 interface and VDSL2 interface on VDSL2 Mini-PIMs. On VDSL2 Mini-PIMs, the `pt-1/0/0` interface is created by default. You can switch to ADSL mode by configuring `at-1/0/0`. You can deactivate `pt-1/0/0` before you create `at-1/0/0` or deactivate `at-1/0/0` to create `pt-1/0/0`. Make sure that you have deleted the previous configurations on `pt-1/0/0` and `pp0`.

In this example:

1. Begin a new configuration on a VDSL2 Mini-PIM.
2. Deactivate previous interfaces and delete the old configuration.
3. Set the interfaces with the VDSL profile and the Layer 3 configuration for the end-to-end data path.
4. Configure the PPPoE on the `pt-1/0/0` interface with a static IP address or CHAP authentication with unnumbered IP address (PAP authentication or CHAP authentication).
5. Configure PPPoE on the `pt-1/0/0` interface with negotiated IP address (PAP authentication or CHAP authentication).

To configure VDSL2 Interfaces in ADSL mode:

1. Configure the ADSL interface for end-to-end data path.
2. Configure PPPoA on the `at-1/0/0` interface with a negotiated IP address and either PAP authentication or CHAP authentication.
3. Configure a static IP address and an unnumbered IP address (and either PAP authentication or CHAP authentication) for PPPoA on the `at-1/0/0` interface.
4. Configure PPPoE on the `at-1/0/0` interface with a negotiated IP address and either PAP authentication or CHAP authentication.

Table 21 on page 130 specifies the CLI quick configuration commands used for configuring VDSL2 interfaces.

Table 21: CLI Quick Configuration

Configuration Step	CLI Quick Configuration Commands
Configure the VDSL2 interface and enable VLAN tagging	**set interfaces pt-1/0/0 vdsl-options vdsl-profile auto** **set interfaces pt-1/0/0 vlan-tagging** **set interfaces pt-1/0/0 unit 0 vlan-id 100**
Begin a new configuration on a VDSL2 Mini-PIM	[edit] **deactivate interface pt-1/0/0** **deactivate interface at-1/0/0** **delete interface pt-1/0/0** **delete interface pp0**
Configure VDSL2 Mini-PIM for End-to-End Data Path	**set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a** **set interfaces pt-1/0/0 unit 0 family inet address 11.11.11.1/24**
Configure PPPoE on the pt-1/0/0 Interface with a Static IP Address	user@host# **set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a** user@host# **set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether** user@host# **set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name locky local-password india passive** user@host# **set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0 auto-reconnect 120 client** user@host# **set interfaces pp0 unit 0 family inet address 10.1.1.6/24** user@host# **set access profile pap_prof authentication-order password client cuttack pap-password india**

Table 21: CLI Quick Configuration *(Continued)*

Configuration Step	CLI Quick Configuration Commands
Configure PPPoE on the pt- Interface with a Static IP Address (CHAP Authentication)	user@host# **set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a** user@host# **set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether** user@host# **set interfaces pp0 unit 0 ppp-options chap default-chap-secret india local-name locky passive** user@host# **set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0 auto-reconnect 120 client** user@host# **set interfaces pp0 unit 0 family inet address 10.1.1.6/24**
Configure PPPoE on the pt-x/x/x Interface with Unnumbered IP (PAP Authentication)	user@host# **set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a** user@host# **set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether** user@host# **set interfaces lo0 unit 0 family inet address 10.1.1.24/32** user@host# **set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof local-name locky local-password india passive** user@host# **set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0 auto-reconnect 120 client** user@host# **set interfaces pp0 unit 0 family inet unnumbered-address lo0.0 destination 10.1.1.1** user@host# **set access profile pap_prof authentication-order password client cuttack pap-password india**

Table 21: CLI Quick Configuration *(Continued)*

Configuration Step	CLI Quick Configuration Commands
Configure PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication)	user@host# **set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a** user@host# **set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether** user@host# **set interfaces lo0 unit 0 family inet address 10.1.1.24/32** user@host# **set interfaces pp0 unit 0 ppp-options chap default-chap-secret india local-name locky passive** user@host# **set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0 auto-reconnect 120 client** user@host# **set interfaces pp0 unit 0 family inet unnumbered-address lo0.0 destination 10.1.1.1**
Configure PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication)	user@host# **set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a** user@host# **set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether** user@host# **set interfaces pp0 unit 0 ppp-options pap access-profile my_prf local-name purple local-password <password> passive** user@host# **set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0 auto-reconnect 120 client** user@host# **set interfaces pp0 unit 0 family inet negotiate-address** user@host# **set access profile my_prf authentication-order password** user@host# **set access profile my_prf**

Table 21: CLI Quick Configuration *(Continued)*

Configuration Step	CLI Quick Configuration Commands
Configure PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication)	user@host# **set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a** user@host# **set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether** user@host# **set interfaces pp0 unit 0 ppp-options chap default-chap-secret <password> local-name purple passive** user@host# **set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0 auto-reconnect 120 client** user@host# **set interfaces pp0 unit 0 family inet negotiate-address**

Configure the VDSL2 Interface and Enable VLAN Tagging

In this example, you create a VDSL2 interface called pt-1/0/0 and set the VDSL2 profile to auto. For more information on basic connectivity refer to Quick Start Guide and to configure network interfaces refer to "Example: Configure Ethernet Interface" on page 166. To configure the VDSL2 interfaces and enable VLAN tagging:

1. Create an interface.

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

2. Set the VDSL2 profile type.

```
[edit interfaces pt-1/0/0]
user@host# set vdsl-options vdsl-profile auto
```

3. Specify the logical unit to connect to the physical VDSL2 interface.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0
```

4. Specify the family protocol type.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0 family inet address 100.100.100.1/24
```

5. Enable VLAN tagging on the pt- interface.

```
[edit interfaces pt-1/0/0]
user@host# set interface pt-1/0/0 vlan-tagging
```

6. Specify the VLAN ID value.

```
[edit interfaces pt-1/0/0]
user@host# set interface pt-1/0/0 unit 0 vlan-id 100
```

7. Commit the configuration.

   VDSL2 is supported only on the pt- interface.The range of VLANs that can be configured is 0 to 4093.

Similarly, you can configure the VDSL2 interface on Annex B (integrated VDSL2 interfaces in ADSL backward compatible mode). After completing the configuration successfully, view the parameters by using the `show interfaces pt-1/0/0` command.

Configure VDSL2 Interface with VDSL2 Mini-PIMs

This example uses VDSL2 Mini-PIMs. Figure 14 on page 135 shows typical SRX Series devices with VDSL2 Mini-PIM network connections.

Figure 14: SRX Series Device with VDSL2 Mini-PIMs in an End-to-End Deployment Scenario

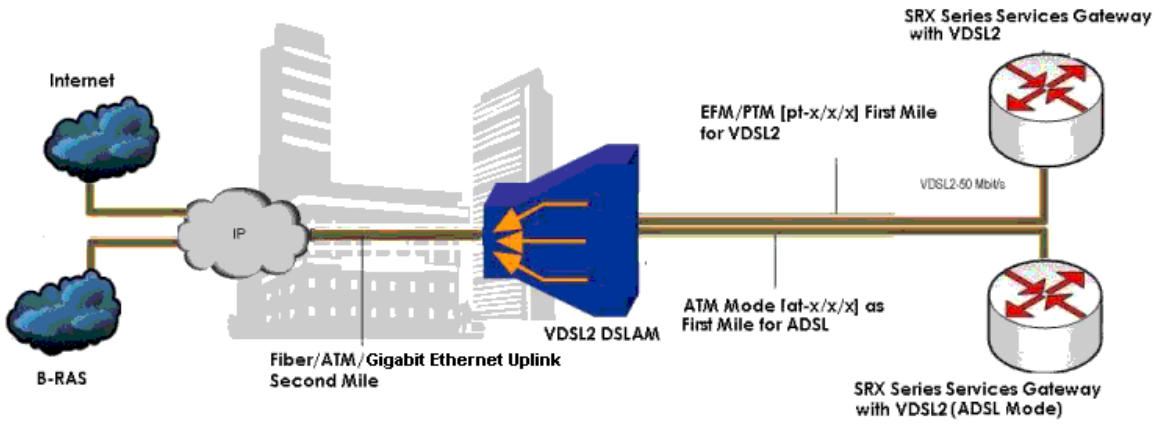

To view the CLI quick configuration commands, see Table 21 on page 130.

To begin a new configuration on a VDSL2 Mini-PIM:

1. Deactivate any previous interfaces.

```
[edit]
user@host# deactivate interface pt-1/0/0
user@host# deactivate interface at-1/0/0
```

2. Delete any old configurations.

```
[edit]
user@host# delete interface pt-1/0/0
user@host# delete interface pp0
```

3. Commit the configuration.

Use the `show chassis fpc` command to see the output of the configuration.

Configure the VDSL2 Mini-PIM for End-to-End Data Path

To configure the VDSL2 Mini-PIM for end-to-end data path:

1. Configure the interfaces with the VDSL profile and the Layer 3 configuration for the end-to-end data path.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 family inet address 11.11.11.1/24
```

2. Commit the configuration.

Use the `show interfaces pt-1/0/0` command to see the output of the configuration.

Configure PPPoE on the pt-1/0/0 Interface with a Static IP Address

To configure the PPPoE on the pt-1/0/0 interface with a static IP address:

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof
user@host# set interfaces pp0 unit 0 ppp-options pap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options pap local-password india
user@host# set interfaces pp0 unit 0 ppp-options pap passive
```

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
```

5. Configure the access profile for the interface.

```
[edit]
user@host# set access profile pap_prof authentication-order password
user@host# set access profile pap_prof client cuttack pap-password india
```

6. Commit the configuration.

Use the `show interfaces pp0`, `show interfaces pt-1/0/0` and `show access profile pap_prof` commands to see the output of the configuration.

Configure PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication)

To configure the PPPoE on the pt-1/0/0 interface with a static IP address (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
user@host# set interfaces pp0 unit 0 ppp-options chap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options chap passive
```

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet address 10.1.1.6/24
```

5. Commit the configuration.

Use the `show interfaces pt-1/0/0` and `show interfaces pp0` commands to see the output of the configuration.

Configure PPPoE on the pt-x/x/x Interface with Unnumbered IP (PAP Authentication)

To configure PPPoE on the pt-1/0/0 interface with unnumbered IP (PAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
```

3. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile pap_prof
user@host# set interfaces pp0 unit 0 ppp-options pap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options pap local-password india
user@host# set interfaces pp0 unit 0 ppp-options pap passive
```

4. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

5. Configure the unnumbered address and destination for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces pp0 unit 0 family inet unnumbered-address destination 10.1.1.1
```

6. Configure the access profile for the interface.

```
[edit]
user@host# set access profile pap_prof authentication-order password
user@host# set access profile pap_prof client cuttack pap-password india
```

7. Commit the configuration.

Use the `show interfaces lo0`, `show interfaces pt-1/0/0`, and `show interfaces pp0` commands to see the output of the configuration.

Configuring PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication)

To configure PPPoE on the pt-1/0/0 interface with unnumbered IP (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the IP address for the interface.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet address 10.1.1.24/32
```

3. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret india
user@host# set interfaces pp0 unit 0 ppp-options chap local-name locky
user@host# set interfaces pp0 unit 0 ppp-options chap passive
```

4. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

5. Configure the unnumbered address and destination for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces pp0 unit 0 family inet unnumbered-address destination 10.1.1.1
```

6. Commit the configuration.

Use the `show interfaces pp0`, `show interfaces pt-1/0/0`, and `show interfaces lo0` commands to see the output of the configuration.

Configure PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication)

To configure PPPoE on the pt-1/0/0 interface with negotiated IP (PAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options pap access-profile my_prf
user@host# set interfaces pp0 unit 0 ppp-options pap local-name purple
user@host# set interfaces pp0 unit 0 ppp-options pap local-password <password>
user@host# set interfaces pp0 unit 0 ppp-options pap passive
```

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
```

```
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the negotiated IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet negotiate-address
```

5. Configure the access profile for the interface.

```
[edit]
user@host# set access profile my_prf authentication-order password
user@host# set access profile my_prf
```

6. Commit the configuration.

Use the `show interfaces pt-1/0/0`, `show interfaces pp0`, and `show access profile my_prf` commands to see the output of the configuration.

Configure PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication)

To configure PPPoE on the pt-1/0/0 interface with negotiated IP (CHAP authentication):

1. Configure the VDSL options and encapsulation for the interface.

```
[edit]
user@host# set interfaces pt-1/0/0 vdsl-options vdsl-profile 17a
user@host# set interfaces pt-1/0/0 unit 0 encapsulation ppp-over-ether
```

2. Configure the PPP options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 ppp-options chap default-chap-secret <password>
user@host# set interfaces pp0 unit 0 ppp-options chap local-name purple
user@host# set interfaces pp0 unit 0 ppp-options chap passive
```

3. Configure the PPPoE options for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 pppoe-options underlying-interface pt-1/0/0.0
user@host# set interfaces pp0 unit 0 pppoe-options auto-reconnect 120
user@host# set interfaces pp0 unit 0 pppoe-options client
```

4. Configure the negotiated IP address for the interface.

```
[edit]
user@host# set interfaces pp0 unit 0 family inet negotiate-address
```

5. Commit the configuration.

Use the `show interfaces pp0` and `show interfaces pt-1/0/0` commands to see the output of the configuration. Similarly, you can configure the integrated VDSL2 interfaces, Annex B, in ADSL backward compatible mode by using the `show interfaces pt-1/0/0` command.

Verification

IN THIS SECTION

- Purpose | 142
- Action | 143

Purpose

Display information about the parameters configured on the VDSL2 interface.

Action

- To display information about the parameters configured on VDSL2 Interface connected to the DSLAM, operating in Annex A and display details of VLAN tagging:

```
user@host>  show interfaces pt-1/0/0
```

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
 Interface index: 146, SNMP ifIndex: 524, Generation: 149
 Type: PTM, Link-level type: Ethernet, MTU: 1496, VDSL mode, Speed:    45440kbps
 Speed: VDSL2
 Device flags : Present Running
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:b1:7e:85:84:ff
 Last flapped : 2009-10-18 11:56:50 PDT (12:32:49 ago)
 Statistics last cleared: 2009-10-19 00:29:37 PDT (00:00:02 ago)
 Traffic statistics:
  Input bytes : 22438962 97070256 bps
  Output bytes : 10866024 43334088 bps
  Input packets: 15141 8187 pps
  Output packets: 7332 3655 pps
Input errors:
   Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0,
   L2 channel errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
 Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 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 6759 6760 0
    1 expedited-fo 0 0 0
    2 assured-forw 0 0 0
    3 network-cont 0 0 0
VDSL alarms : None
VDSL defects : None
VDSL media: Seconds Count State
   LOF 0 0 OK
   LOS 0 0 OK
   LOM 0 0 OK
```

```
   LOP 0 0 OK
   LOCDI 0 0 OK
   LOCDNI 0 0 OK
VDSL status:
  Modem status : Showtime (Profile-17a)
  VDSL profile : Profile-17a Annex A
  Last fail code: None
  Subfunction : 0x00
  Seconds in showtime : 45171
continue.....................................................
.....................................................
```

Similarly, you can verify the VDSL2 interface on Annex B mode by using the `show interfaces pt-1/0/0` command.

```
user@host>  show interfaces pt-1/0/0
```

```
vlan-tagging;
vdsl-options {
vdsl-profile auto;
}
unit 0 {
vlan-id 100;
Family inet {
address 100.100.100.1/24;
}
```

- Verify the FPC status by entering the `show chassis fpc` command. The VDSL2 Mini-PIM is installed in the first slot of the SRX320 device chassis; therefore, use `fpc 1`. For SRX340 devices, use the FPCs `fpc 1`, `fpc 2`, `fpc 3`, or `fpc 4`.

```
user@host>  show chassis fpc
```

```
Temp CPU Utilization (%) Memory Utilization
(%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
```

```
0 Online -------------------- CPU less FPC --------------------
1 Online -------------------- CPU less FPC --------------------
```

- Verify the status of interface, modem status, time in seconds and VDSL profile of DSLAM by using the `run show interface pt-1/0/0`.

```
user@host> show interface pt-1/0/0
```

```
Physical interface: pt-1/0/0, Enabled, Physical link is Up
 Interface index: 146, SNMP ifIndex: 524, Generation: 149
 Type: PTM, Link-level type: Ethernet, MTU: 1496, VDSL mode, Speed:    45440kbps
 Speed: VDSL2
 Device flags : Present Running
 Link flags : None
 CoS queues : 8 supported, 8 maximum usable queues
 Hold-times : Up 0 ms, Down 0 ms
 Current address: 00:b1:7e:85:84:ff
 Last flapped : 2009-10-18 11:56:50 PDT (12:32:49 ago)
 Statistics last cleared: 2009-10-19 00:29:37 PDT (00:00:02 ago)
 Traffic statistics:
  Input bytes : 22438962 97070256 bps
  Output bytes : 10866024 43334088 bps
  Input packets: 15141 8187 pps
  Output packets: 7332 3655 pps
Input errors:
   Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0,
   L2 channel errors: 0, L2 mismatch timeouts: 0, Resource errors: 0
 Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 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 6759 6760 0
    1 expedited-fo 0 0 0
    2 assured-forw 0 0 0
    3 network-cont 0 0 0
VDSL alarms : None
VDSL defects : None
VDSL media: Seconds Count State
   LOF 0 0 OK
```

```
   LOS 0 0 OK
   LOM 0 0 OK
   LOP 0 0 OK
   LOCDI 0 0 OK
   LOCDNI 0 0 OK
VDSL status:
  Modem status : Showtime (Profile-17a)
  VDSL profile : Profile-17a Annex A
  Last fail code: None
  Subfunction : 0x00
  Seconds in showtime : 45171
continue.....................................................
.............................................
```

- To display all the parameters configured on VDSL2 Mini-PIM for End-to-End Data Path

```
user@host> show interfaces pt-1/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up   inet     11.11.11.1/24

[edit]
user@host# run ping 11.11.11.2 count 1000 rapid
PING 11.11.11.2 (11.11.11.2): 56 data bytes
!!!
```

```
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:b1:7e:85:84:ff
  Last flapped   : 2009-10-28 00:36:29 PDT (00:12:03 ago)
  Statistics last cleared: 2009-10-28 00:47:56 PDT (00:00:36 ago)
  Traffic statistics:
   Input  bytes  :                84000                    0 bps
   Output bytes  :               138000                    0 bps
   Input  packets:                 1000                    0 pps
   Output packets:                 1000                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0, L2
mismatch timeouts: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 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                1000                 1000                    0
    1 expedited-fo                  0                    0                    0
    2 assured-forw                  0                    0                    0
    3 network-cont                  0                    0                    0
  VDSL  alarms   : None
  VDSL  defects  : None
  VDSL media:            Seconds        Count  State
    LOF                          0            0  OK
    LOS                          0            0  OK
    LOM                          0            0  OK
    LOP                          0            0  OK
    LOCDI                        0            0  OK
    LOCDNI                       0            0  OK
  VDSL status:
    Modem status  : Showtime (Profile-17a)
    VDSL profile  : Profile-17a Annex A
    Last fail code: None
```

- To display the PPPoE on the pt-1/0/0 Interface with a Static IP Address.

```
user@host> show interfaces pp0
```

```
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 510
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
    Input packets : 0
    Output packets: 0

  Logical interface pp0.0 (Index 71) (SNMP ifIndex 522)
    Flags: Hardware-Down Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionDown, Session ID: None,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 57
    Output packets: 56
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 22 (00:00:40 ago), Output: 25 (00:00:04 ago)
  LCP state: Down
```

```
user@host> show interfaces pt-1/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up

[edit]
user@host# run show interfaces pp0 terse
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.6/24
```

```
[edit]
user@host# run ping 10.1.1.1 count 100 rapid
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!
```

```
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.6
```

```
user@host> show interfaces pt-1/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up
```

```
user@host> show interfaces pp0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.6/24
```

```
user@host> ping 10.1.1.1 count 100 rapid
```

```
PING 10.1.1.1 (10.1.1.1): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!
--- 10.1.1.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.608/15.466/25.939/1.779 ms
```

- To display the PPPoE on the pt-1/0/0 Interface with Unnumbered IP (PAP Authentication)

```
user@host> show interfaces pp0
```

```
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 510
  Type: PPPoE, Link-level type: PPPoE, MTU: 1532
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
    Input packets : 0
    Output packets: 0

  Logical interface pp0.0 (Index 72) (SNMP ifIndex 522)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 33,
      Session AC name: cuttack, Remote MAC address: 00:03:6c:c8:8c:55,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 22
    Output packets: 20
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1 (00:00:08 ago), Output: 0 (never)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
  CHAP state: Closed
  PAP state: Success
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1.1, Local: 10.1.1.24
```

```
user@host> show interfaces pt-1/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up
```

```
user@host> show interfaces pp0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.24           --> 10.1.1.1
```

```
user@host> ping 10.1.1.1 count 100 rapid
```

```
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!
```

```
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
    Input packets : 0
    Output packets: 0

  Logical interface pp0.0 (Index 70) (SNMP ifIndex 522)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 35,
      Session AC name: cuttack, Remote MAC address: 00:03:6c:c8:8c:55,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
      Underlying interface: pt-1/0/0.0 (Index 69)
    Input packets : 25
    Output packets: 22
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 2 (00:00:10 ago), Output: 2 (00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
  CHAP state: Success
  PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1.1, Local: 10.1.1.24
```

```
user@host> show interfaces pt-1/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up
```

```
user@host> show interfaces pp0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     10.1.1.24           --> 10.1.1.1
```

```
user@host> ping 10.1.1.1 count 100 rapid
```

```
PING 10.1.1.1 (10.1.1.1): 56 data bytes
!!!
```

```
Interface flags: Point-To-Point SNMP-Traps
Link type      : Full-Duplex
Link flags     : None
  Input packets : 0
  Output packets: 0

Logical interface pp0.0 (Index 72) (SNMP ifIndex 522)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 4,
    Session AC name: belur, Remote MAC address: 00:90:1a:43:18:d1,
    Configured AC name: None, Service name: None,
    Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
    Underlying interface: pt-1/0/0.0 (Index 69)
  Input packets : 18
  Output packets: 18
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 0 (never), Output: 11 (00:00:01 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
CHAP state: Closed
PAP state: Success
  Security: Zone: Null
  Protocol inet, MTU: 1474
    Flags: Negotiate-Address
```

```
      Addresses, Flags: Kernel Is-Preferred Is-Primary
        Destination: 12.12.12.1, Local: 12.12.12.11
```

```
user@host> show interfaces pt-1/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up
```

```
user@host> show interfaces pp0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     12.12.12.11         --> 12.12.12.1
```

```
user@host> ping 12.12.12.1 count 100 rapid
```

```
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!
```

```
Interface flags: Point-To-Point SNMP-Traps
Link type      : Full-Duplex
Link flags     : None
  Input packets : 0
  Output packets: 0

Logical interface pp0.0 (Index 70) (SNMP ifIndex 522)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 8,
    Session AC name: belur, Remote MAC address: 00:90:1a:43:18:d1,
    Configured AC name: None, Service name: None,
    Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
    Underlying interface: pt-1/0/0.0 (Index 69)
  Input packets : 12
  Output packets: 11
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 0 (never), Output: 4 (00:00:03 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
CHAP state: Success
PAP state: Closed
  Security: Zone: Null
  Protocol inet, MTU: 1474
    Flags: Negotiate-Address
```

```
      Addresses, Flags: Kernel Is-Preferred Is-Primary
        Destination: 12.12.12.1, Local: 12.12.12.12
```

```
user@host> show interfaces pt-1/0/0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pt-1/0/0                up    up
pt-1/0/0.0              up    up
```

```
user@host> show interfaces pp0 terse
```

```
Interface               Admin Link Proto    Local                 Remote
pp0                     up    up
pp0.0                   up    up   inet     12.12.12.12         --> 12.12.12.1
```

```
user@host> ping 12.12.12.1 count 100 rapid
```

```
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!
```

4

CHAPTER

Configuring Ethernet Interfaces

Ethernet Interfaces | 160

Configuring Aggregated Ethernet Interfaces | 171

Configuring Link Aggregation Control Protocol | 188

Configuring Gigabit Ethernet Physical Interface Modules | 207

Port Speed on SRX Series Devices | 251

Power over Ethernet | 261

Ethernet Interfaces

IN THIS SECTION

- Ethernet Interfaces Overview | 160
- Example: Configure Ethernet Interface | 166
- Example: Configuring Promiscuous Mode on the SRX5K-MPC | 167

Learn about Ethernet technology used to broadcast traffic on security devices, static ARP entries, creating and deleting the Ethernet interface, and enabling and disabling the promiscuous mode on these interfaces. Also learn about Aggregated Ethernet Interfaces

Ethernet Interfaces Overview

IN THIS SECTION

- Ethernet Frames | 164
- Promiscuous Mode | 165

Ethernet is a Layer 2, point-to multipoint technology that operates in a shared bus topology, supports broadcast transmission, and has distributed access control.

In a shared bus topology, all devices connect to a single, shared physical link through which all data transmissions are sent. The devices within a single Ethernet topology make up a broadcast domain.

The physical hardware does not provide information to the sender about incoming and lost traffic. Higher layer protocols such as TCP/IP can provide this type of notification.

Table 22: Types of Ethernet Interfaces

Types	Description
Ethernet Access Control and Transmission	• Ethernet's access control is distributed. • Uses carrier-sense multiple access with collision detection (CSMA/CD) mechanism. • If there is no transmission host begins transmitting its own data. • Length of each transmission is determined by fixed Ethernet packet size. • Enforces a minimum idle time between transmissions. • Ensures there is no interruption in sending and receiving traffic.

Collisions and Detection	• Delay, or latency, in transmitting traffic results in collision of two electrical signals.
	• Signals are scrambled so that both transmissions are effectively lost .
	• Two types include: Collision detection and Backoff Algorithm
	  • Collision detection refers to link monitoring while the devices are transmitting data. The device transmits data during the idle state on the wire.
	  • Binary exponential backoff algorithm helps each device, sending a colliding transmission randomly, select a value within a range. The value represents the number of transmission times that the device must wait before retransmitting its data. Each time a collision occurs, the range of values doubles.

Table 22: Types of Ethernet Interfaces *(Continued)*

Types	Description

Collision Domains and LAN Segments	• Multiple collision domains can be interconnected by repeaters, bridges, and switches if the length of an Ethernet cable restrict the length of a LAN segment.
	• Repeaters are electronic devices that act on analog signals and relay all electronic signals. Ethernet specification restricts the number of repeaters to two. A single repeater can double the distance between two devices on an Ethernet network.
	• Bridges and switches combine LAN segments into a single Ethernet network by using multiple ports to connect the physical wires in each segment.
	• Bridges provide more management and interface ports.
	• Bridge tracks the source MAC address of the packets and stores the addresses and their associated input ports in an interface table.
	• The bridge examines its interface table and takes one of the following actions:
	  • If the destination address does not match an interface table address, the bridge transmits the packet to all hosts on the network using the Ethernet broadcast address.
	  • If the destination address matches the port with receiving packet, the bridge or switch discards the packet. The bridge does not need to retransmit it.
	  • If the destination address maps to a port other than the one through which the packet was received, the bridge transmits the packet through the appropriate port to the corresponding LAN segment.

	• Combination of all the LAN segments within an Ethernet network is called broadcast domain. • When you use a bridge or switch, the broadcast domain consists of the entire LAN.

Table 23 on page 163

Table 23: Collision Backoff Algorithm Rounds

Round	Size of Set	Elements in the Set
1	2	{0,1}
2	4	{0,1,2,3}
3	8	{0,1,2,3,...,7}
4	16	{0,1,2,3,4,...,15}
5	32	{0,1,2,3,4,5,...,31}
6	64	{0,1,2,3,4,5,6,...,63}
7	128	{0,1,2,3,4,5,6,7,...,127}
8	256	{0,1,2,3,4,5,6,7,8,...,255}
9	512	{0,1,2,3,4,5,6,7,8,9,...,511}
10	1024	{0,1,2,3,4,5,6,7,8,9,10,...,1023}

Ethernet Frames

Data is transmitted through an Ethernet network in frames. The frames are of variable length, ranging from 64 octets to 1518 octets, including the header, payload, and cyclic redundancy check (CRC) value.Figure 15 on page 164 shows the Ethernet frame format.

Figure 15: Ethernet Frame Format

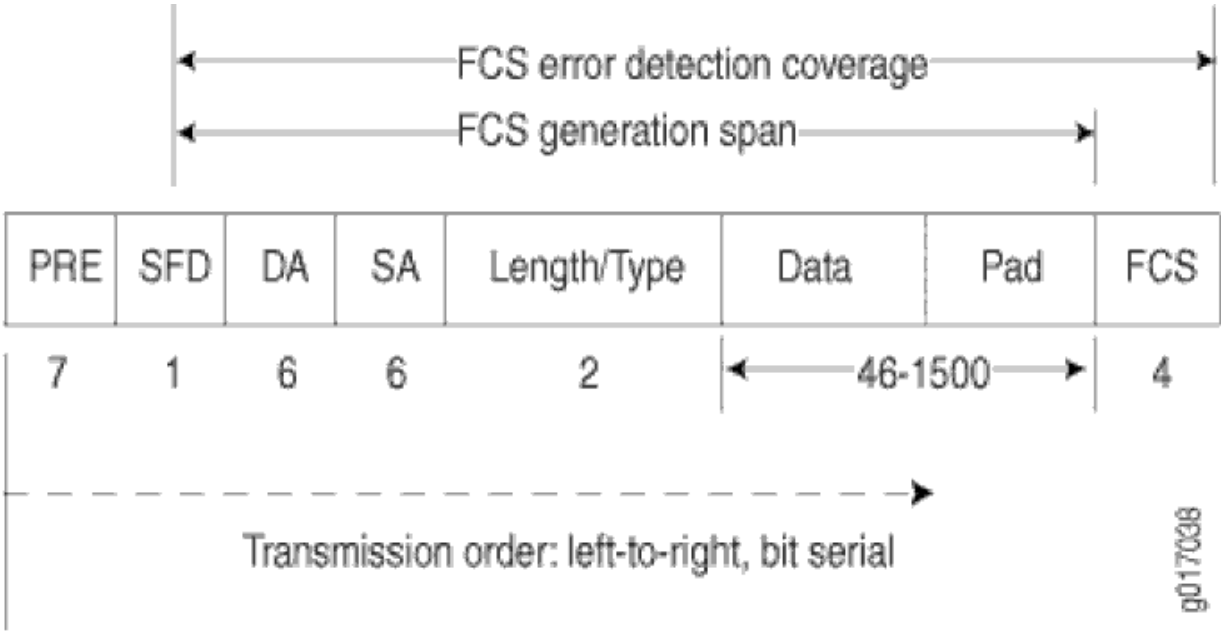

Ethernet frames have the following fields:

- The preamble (PRE) field is 7 octets of alternating 0s and 1s. The predictable format in the preamble allows receiving interfaces to synchronize themselves to the data being sent. The preamble is followed by a 1-octet start-of-frame delimiter (SFD).
- The destination address (DA) and source address (SA) fields contain the 6-octet (48-bit) MAC addresses for the destination and source ports on the network. These Layer 2 addresses uniquely identify the devices on the LAN.
- The Length/Type field is a 2-octet field that either indicates the length of the frame's data field or identifies the protocol stack associated with the frame. Here are some common frame types:
  - AppleTalk—`0x809B`
  - AppleTalk ARP—`0x80F3`
  - DECnet—`0x6003`

  - IP—`0x0800`
  - IPX—`0x8137`
  - Loopback—`0x9000`
  - XNS—`0x0600`
- The Data field contains the packet payload.
- The frame check sequence (FCS) is a 4-octet field that contains the calculated CRC value. This value is calculated by the originating host and appended to the frame. When it receives the frames, the receiving host calculates the CRC and checks it against this appended value to verify the integrity of the received frame.
- On SRX650 devices, MAC pause frame and FCS error frame counters are not supported for the interfaces ge-0/0/0 through ge-0/0/3. (Platform support depends on the Junos OS Release in your installation.)

Promiscuous Mode

- When you enable promiscuous mode on a Layer 3 Ethernet interface, all received packets on are sent to the central point or Services Processing Unit (SPU) regardless of the destination MAC address of the packet.
- You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and aggregated Ethernet interfaces.
- If you enable promiscuous mode on a redundant Ethernet interface, it is enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, it is enabled on all member interfaces.
- Promiscuous mode function is supported on 1-Gigabit, 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces on the I/O cards (IOCs) and the SRX5000 line Module Port Concentrator (SRX5K-MPC).
- By default, an interface enables MAC filtering. You can configure promiscuous mode on the interface to disable MAC filtering. When you delete the configuration, the interface will perform MAC filtering again.
- You can change the interface MAC address when the interface is operating in promiscuous mode. When the interface is operating in normal mode, the MAC filtering function on the IOC uses the new MAC address to filter the packets.

Example: Configure Ethernet Interface

IN THIS SECTION

- Overview | 166

Overview

Table describes the steps to create and (optional) delete Ethernet interfaces on your routing device.

Table 24: Ethernet Interfaces Configuration

Configuration Step	Command
Step 1: Create the Ethernet interface and set the logical interface.	[edit] user@host# **edit interfaces ge-1/0/0 unit 0**
Step 2: If you are done configuring the device, commit the configuration.	[edit] user@host# **commit**
Step 3: (Optional) Specify the interface you want to delete.	[edit] user@host# **delete interfaces ge-1/0/0**
Step 4: If you are done configuring the device, commit the configuration.	[edit] user@host# **commit**

Example: Configuring Promiscuous Mode on the SRX5K-MPC

IN THIS SECTION

- Verification | 168

This example shows how to configure promiscuous mode on an SRX5K-MPC interface in an SRX5600 to disable MAC address filtering.

CLI Quick Configuration

Below table specifies the CLI quick configuration commands used for configuring and disabling promiscuous mode on SRX5K-MPC interface .

Table 25: CLI Quick Configuration

Configuration Step	CLI Quick Configuration Commands
Configure promiscuous mode on the interface	**set interfaces et-4/0/0 unit 0 family inet address 10.1.1.1/24** **set interfaces et-4/0/0 promiscuous-mode**
Disable promiscuous mode on an interface	user@host# **delete interfaces et-4/0/0 promiscuous-mode**

Configure Promiscuous Mode on an Interface

Below table describes the step-by-step to configure promiscuous mode on an interface on your security device.

Table 26: Promiscuous Mode Configuration

Configuration Step	Command
Step 1: Configure the ingress interface.	[edit interfaces] user@host# **set et-4/0/0 unit 0 family inet address 10.1.1.1/24**
Step 2: Enable promiscuous mode on the interface.	[edit interfaces] user@host# **set et-4/0/0 promiscuous-mode**
Step 3: (Optional) Disable promiscuous mode on the interface.	[edit] user@host# **delete interfaces et-4/0/0 promiscuous-mode**

Use the `show interfaces` command to see the output of the configuration.

Verification

Purpose

Verify that promiscuous mode is enabled, its status, on the interface and disabled on the interface.

Action

- To display information about the parameters configured on promiscuous mode Interface.

```
user@host> show interfaces
```

```
Physical interface: et-4/0/0, Enabled, Physical link is Up
  Interface index: 137, SNMP ifIndex: 511
  Link-level type: Ethernet, MTU: 1518, Speed: 100Gbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: Promiscuous SNMP-Traps Internal: 0x4000
  CoS queues     : 8 supported, 8 maximum usable queues
```

```
  Current address: 2c:21:72:3a:05:28, Hardware address: 2c:21:72:3a:05:28
  Last flapped   : 2014-01-17 14:44:53 PST (5d 06:30 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0

  Logical interface et-4/0/0.0 (Index 71) (SNMP ifIndex 513)
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1351 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: HOST
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
    ospf pgm pim rip router-discovery rsvp sap vrrp
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 122.122.122/24, Local: 122.122.122.1,
        Broadcast: 122.122.122.255
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary

  Logical interface et-4/0/0.32767 (Index 72) (SNMP ifIndex 517)
    Flags: SNMP-Traps 0x4004000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: HOST
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
    ospf pgm pim rip router-discovery rsvp sap vrrp
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

The `Interface flags: Promiscuous` field shows that promiscuous mode is enabled on the interface.

- Verify that promiscuous mode works on the `et-4/0/0` interface. Send traffic into the `et-4/0/0` interface with a MAC address that is different from the interface MAC address and turn on promiscuous mode. From operational mode, enter the `monitor interface traffic` command.

```
user@host> monitor interface traffic
```

```
Interface    Link  Input packets       (pps)     Output packets      (pps)
gr-0/0/0      Up              0         (0)               0         (0)
ip-0/0/0      Up              0         (0)               0         (0)
lt-0/0/0      Up              0         (0)               0         (0)
xe-1/2/0    Down              0         (0)               0         (0)
xe-1/2/1    Down              0         (0)               0         (0)
xe-1/2/2    Down              0         (0)               0         (0)
xe-1/2/3    Down              0         (0)               0         (0)
xe-1/2/4    Down              0         (0)               0         (0)
xe-1/2/5    Down              0         (0)               0         (0)
xe-1/2/6    Down              0         (0)               0         (0)
xe-1/2/7    Down              0         (0)               0         (0)
xe-1/2/8    Down              0         (0)               0         (0)
xe-1/2/9    Down              0         (0)               0         (0)
et-4/0/0      Up        4403996     (100002)              0         (0)
et-4/2/0      Up              3         (0)         4403924      (99997)
avs0          Up              0         (0)               0         (0)
avs1          Up              0         (0)               0         (0)
dsc           Up              0                           0
em0           Up          15965                       14056
```

The `input packets` and `pps` fields show that traffic is passing through the `et-4/0/0` interface as expected after promiscuous mode is enabled.

- Verify that disabled promiscuous mode works on the `et-4/0/0` interface. Send traffic and turn off the promiscuous mode.

```
user@host> monitor interface traffic
```

```
Interface    Link  Input packets       (pps)     Output packets      (pps)
gr-0/0/0      Up              0         (0)               0         (0)
ip-0/0/0      Up              0         (0)               0         (0)
```

```
lt-0/0/0      Up               0          (0)               0          (0)
xe-1/2/0      Down             0          (0)               0          (0)
xe-1/2/1      Down             0          (0)               0          (0)
xe-1/2/2      Down             0          (0)               0          (0)
xe-1/2/3      Down             0          (0)               0          (0)
xe-1/2/4      Down             0          (0)               0          (0)
xe-1/2/5      Down             0          (0)               0          (0)
xe-1/2/6      Down             0          (0)               0          (0)
xe-1/2/7      Down             0          (0)               0          (0)
xe-1/2/8      Down             0          (0)               0          (0)
xe-1/2/9      Down             0          (0)               0          (0)
et-4/0/0      Up        11505495          (0)               0          (0)
et-4/2/0      Up               6          (0)        11505425          (0)
avs0          Up               0          (0)               0          (0)
avs1          Up               0          (0)               0          (0)
dsc           Up               0                            0
em0           Up           37964                        31739
```

The pps field shows that the traffic is not passing through the et-4/0/0 interface after promiscuous mode is disabled.

RELATED DOCUMENTATION

- Understanding Interfaces

Configuring Aggregated Ethernet Interfaces

IN THIS SECTION

- Understanding Aggregated Ethernet Interfaces | 172
- Configuring Aggregated Ethernet Interfaces | 174
- Understanding Physical Interfaces for Aggregated Ethernet Interfaces | 175
- Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces | 176
- Understanding Aggregated Ethernet Interface Link Speed | 177
- Example: Configuring Aggregated Ethernet Link Speed | 178

- Understanding Minimum Links for Aggregated Ethernet Interfaces | **179**
- Example: Configuring Aggregated Ethernet Minimum Links | **180**
- Deleting Aggregated Ethernet Interface | **181**
- Example: Deleting Aggregated Ethernet Interfaces | **182**
- Example: Deleting Aggregated Ethernet Interface Contents | **183**
- Understanding VLAN Tagging for Aggregated Ethernet Interfaces | **185**
- Understanding Promiscuous Mode for Aggregated Ethernet Interfaces | **185**
- Verifying Aggregated Ethernet Interfaces | **185**

The below topics discuss the overview Aggregated Ethernet (AE) interfaces on security devices, configuration details of AE interfaces, physical interfaces, AE interface link speed, VLAN tagging for aggregated Ethernet interfaces, and deleting an Aggregated Ethernet interface in security devices.

Understanding Aggregated Ethernet Interfaces

IN THIS SECTION

- LAGs | **173**
- LACP | **173**

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. Junos OS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on Layer 3 information carried in the packet, Layer 4 information carried in the packet, or both, or based on session ID data. (The session ID data has higher precedence than the Layer 3 or 4 information.) This implementation uses the same load-balancing algorithm used for per-packet load balancing.

Aggregated Ethernet interfaces can be Layer 3 interfaces (VLAN-tagged or untagged) and Layer 2 interfaces.

NOTE: This topic is specific to the SRX3000 and SRX5000 line devices. For information about link aggregation for other SRX Series devices, see the Ethernet Switching User Guide.

This topic contains the following sections:

LAGs

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle, such that a media access control (MAC) client can treat the LAG as if it were a single link. Support for LAGs based on IEEE 802.3ad makes it possible to aggregate physical interface links on your device. LAGs provide increased interface bandwidth and link availability by linking physical ports and load-balancing traffic crossing the combined interface. For the LAG to operate correctly, it is necessary to coordinate the two end systems connected by the LAG, either manually or automatically.

Internally, a LAG is a virtual interface presented on SRX3000 and SRX5000 line devices or on any system (consisting of devices such as routers and switches) supporting 802.3ad link aggregation. Externally, a LAG corresponds to a bundle of physical Ethernet links connected between an SRX3000 or SRX5000 line device and another system capable of link aggregation. This bundle of physical links is a virtual link.

Follow these guidelines for aggregated Ethernet support for the SRX3000 and SRX5000 lines:

- The devices support a maximum of 16 physical interfaces per single aggregated Ethernet bundle.
- Aggregated Ethernet interfaces can use interfaces from the same or different Flexible PIC Concentrators (FPCs) and PICs.
- On the aggregated bundle, capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available.

LACP

Junos OS supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs.

Starting with Junos OS Release 15.1X49-D40, LACP is supported on Layer 2 transparent mode in addition to existing support on Layer 3 mode. For information about link aggregation for other SRX Series devices, see the Ethernet Switching User Guide.

LACP provides a standardized means for exchanging information between partner (remote or far-end of the link) systems on a link. This exchange allows their link aggregation control instances to reach

agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG. This exchange also enables the transmission and reception processes for the link to function in an orderly manner.

For example, when LACP is not enabled, a local LAG might attempt to transmit packets to a remote individual interface, which causes the communication to fail. (An individual interface is a nonaggregatable interface.) 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.

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device. Then you associate a set of ports that have the same speed and are in full-duplex mode. The physical ports can be 100-megabit Ethernet, 1-Gigabit Ethernet, and 10-Gigabit Ethernet.

When configuring LACP, follow these guidelines:

- LACP does not support automatic configuration on SRX3000 and SRX5000 line devices, but partner systems are allowed to perform automatic configuration. When an SRX3000 or SRX5000 line device is connected to a fully 802.3ad-compliant partner system, static configuration of LAGs is initiated on the SRX3000 and SRX5000 line device side, and static configuration is not needed on the partner side.
- When an SRX3000 or SRX5000 line device is connected to a Juniper Networks MX Series router, static configuration of LAGs is needed at both the actor (local or near-end of the link) and partner systems.
- Although the LACP functions on the SRX3000 and SRX5000 line devices are similar to the LACP features on Juniper Networks MX Series routers, the following LACP features on MX Series routers are not supported on SRX3000 and SRX5000 line devices: link protection, system priority, and port priority for aggregated Ethernet interfaces. Instead, SRX3000 and SRX5000 line devices provide active/standby support with redundant Ethernet interface LAGs in *chassis cluster* deployments.

LACP is supported in standalone deployments, where aggregated Ethernet interfaces are supported, and in chassis cluster deployments, where aggregated Ethernet interfaces and redundant Ethernet interfaces are supported simultaneously.

Configuring Aggregated Ethernet Interfaces

NOTE: This topic is specific to the SRX3000 and SRX5000 line devices.

To configure an aggregated Ethernet interface:

1. Set the number of aggregated Ethernet interfaces on the device. See Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device.
2. Associate a physical interface with the aggregated Ethernet interface. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 176.
3. (Optional) Set the required link speed for all the interfaces included in the bundle. See "Example: Configuring Aggregated Ethernet Link Speed" on page 178.
4. (Optional) Configure the minimum number of links that must be up for the bundle as a whole to be labeled as up. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 180.
5. (Optional) Enable or disable VLAN tagging. See "Understanding VLAN Tagging for Aggregated Ethernet Interfaces" on page 185.
6. (Optional) Enable promiscuous mode. See "Understanding Promiscuous Mode for Aggregated Ethernet Interfaces" on page 185.

SEE ALSO

Ethernet Switching User Guide

Understanding Physical Interfaces for Aggregated Ethernet Interfaces

You associate a physical interface with an aggregated Ethernet interface. Doing so associates the physical child links with the logical aggregated parent interface to form a link aggregation group (LAG). You must also specify the constituent physical links by including the `802.3ad` *configuration statement*.

A physical interface can be added to any aggregated Ethernet interface as long as all member links have the same link speed and the maximum number of member links does not exceed 16. The aggregated Ethernet interface instance number ae*x* can be from 0 through 127, for a total of 128 aggregated interfaces.

NOTE:

- If you specify (on purpose or accidentally) that a link already associated with an aggregated Ethernet interface be associated with another aggregated Ethernet interface, the link is removed from the previous interface (there is no need for you to explicitly delete it) and it is added to the other one.

- On SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550M devices, when you create an aggregated interface with two or more ports and if a link in the bundle goes down, the traffic forwarded through the same link will be rerouted two seconds later. This causes an outage for the traffic being sent to the link until reroute is complete.

Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces

IN THIS SECTION

- Requirements | 176
- Overview | 176
- Configuration | 176
- Verification | 177

This example shows how to associate physical interfaces with aggregated Ethernet interfaces.

Requirements

Before you begin, set the number of aggregated Ethernet interfaces on the device. See Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device.

Overview

In this example, you associate the physical child link of the ge-1/0/0 and ge-2/0/0 physical interfaces with the logical aggregate parent, ae0, thereby creating a LAG. Similarly, you create a LAG that associate the ge-3/0/0, ge-3/0/1, and ge-4/0/1 physical interfaces with the ae1 aggregated Ethernet interface.

Configuration

IN THIS SECTION

- Procedure | 177

Procedure

Step-by-Step Procedure

To associate physical interfaces with aggregated Ethernet interfaces:

1. Create the first LAG.

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

2. Create the second LAG.

```
[edit]
user@host# set interfaces ge-3/0/0 gigether-options 802.3ad ae1
user@host# set interfaces ge-3/0/1 gigether-options 802.3ad ae1
user@host# sset interfaces ge-4/0/0 gigether-options 802.3ad ae1
```

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces` command.

Understanding Aggregated Ethernet Interface Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. All interfaces that make up a bundle must be the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed you specify in the `link-speed` parameter, an error message will be logged.

The speed value is specified 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 SRX3000 and SRX5000 line devices can have one of the following speed values:

- 100m—Links are 100 Mbps.
- 10g—Links are 10 Gbps.
- 1g—Links are 1 Gbps.

Example: Configuring Aggregated Ethernet Link Speed

IN THIS SECTION

- Requirements | 178
- Overview | 178
- Configuration | 179
- Verification | 179

This example shows how to configure the aggregated Ethernet link speed.

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 176.

Overview

In this example, you set the required link speed for all interfaces included in the bundle to 10 Gbps. All interfaces that make up a bundle must be the same speed.

Configuration

IN THIS SECTION

- Procedure | 179

Procedure

Step-by-Step Procedure

To configure the aggregated Ethernet link speed:

1. Set the link speed.

```
[edit]
user@host# set interfaces ae0 aggregated-ether-options link-speed 10g
```

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces` command.

Understanding Minimum Links for Aggregated Ethernet Interfaces

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

On SRX1000, SRX3000, and SRX5000 line devices, the valid range for the minimum links number is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled as up.

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

Example: Configuring Aggregated Ethernet Minimum Links

IN THIS SECTION

- Requirements | **180**
- Overview | **180**
- Configuration | **181**
- Verification | **181**

This example shows how to configure the minimum number of links on an aggregated Ethernet interface that must be up for the bundle as a whole to be labeled as up.

Requirements

Before you begin:

- Add the aggregated Ethernet interfaces using the device count. See Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 176.
- Configure the aggregated Ethernet link speed. See "Example: Configuring Aggregated Ethernet Link Speed" on page 178.

Overview

In this example, you specify that on interface ae0 at least eight links must be up for the bundle as a whole to be labeled as up.

Configuration

IN THIS SECTION

- Procedure | 181

Procedure

Step-by-Step Procedure

To configure the minimum number of links on an aggregated Ethernet interface:

1. Set the minimum number of links.

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

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces` command.

Deleting Aggregated Ethernet Interface

You can delete an aggregated Ethernet interface from the interface configuration. Junos OS removes the configuration statements related to `ae`*x* and sets this interface to the down state. The deleted aggregated Ethernet interface still exists, but it becomes an empty interface.

Example: Deleting Aggregated Ethernet Interfaces

IN THIS SECTION

- Requirements | 182
- Overview | 182
- Configuration | 182
- Verification | 183

This example shows how to delete aggregated Ethernet interfaces using the device count.

Requirements

Before you begin, set the number of aggregated Ethernet interfaces on the device. See Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device.

Overview

This example shows how to clean up unused aggregated Ethernet interfaces. In this example, you reduce the number of interfaces from 10 to 6, thereby removing the last 4 interfaces from the interface object list.

Configuration

IN THIS SECTION

- Procedure | 182

Procedure

Step-by-Step Procedure

To delete an interface:

1. Set the number of aggregated Ethernet interfaces.

```
[edit]
user@host# delete chassis aggregated-devices ethernet device-count 6
```

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

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

Verification

To verify the configuration is working properly, enter the `show chassis aggregated-devices` command.

Example: Deleting Aggregated Ethernet Interface Contents

IN THIS SECTION

- Requirements | 183
- Overview | 184
- Configuration | 184
- Verification | 184

This example shows how to delete the contents of an aggregated Ethernet interface.

Requirements

Before you begin:

- Set the number of aggregated Ethernet interfaces on the device. See Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device.
- Associate a physical interface with the aggregated Ethernet interface. See "Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces" on page 176.

- Set the required link speed for all the interfaces included in the bundle. See "Example: Configuring Aggregated Ethernet Link Speed" on page 178.
- Configure the minimum number of links that must be up for the bundle as a whole to be labeled as up. See "Example: Configuring Aggregated Ethernet Minimum Links" on page 180.

Overview

In this example, you delete the contents of the ae4 aggregated Ethernet interface, which sets it to the down state.

Configuration

IN THIS SECTION

- Procedure | 184

Procedure

Step-by-Step Procedure

To delete the contents of an aggregated Ethernet interface:

1. Delete the interface.

```
[edit]
user@host# delete interfaces ae4
```

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces` command.

Understanding VLAN Tagging for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can be either VLAN-tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on the SRX3000 and SRX5000 lines support the configuration of `native-vlan-id`, which consists of the following configuration statements:

- `inner-tag-protocol-id`
- `inner-vlan-id`
- `pop-pop`
- `pop-swap`
- `push-push`
- `swap-push`
- `swap-swap`

Understanding Promiscuous Mode for Aggregated Ethernet Interfaces

You can enable promiscuous mode on aggregated Ethernet interfaces. When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or Services Processing Unit (SPU) regardless of the destination MAC address of the packet. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Verifying Aggregated Ethernet Interfaces

IN THIS SECTION

- Verifying Aggregated Ethernet Interfaces (terse) | **186**
- Verifying Aggregated Ethernet Interfaces (extensive) | **187**

Verifying Aggregated Ethernet Interfaces (terse)

IN THIS SECTION

- Purpose | 186
- Action | 186

Purpose

Display status information in terse (concise) format for aggregated Ethernet interfaces.

Action

From operational mode, enter the `show interfaces ae0 terse` command.

```
user@host> show interfaces ae0 terse
ge-2/0/0.0              up    up   aenet    --> ae0.0
ge-2/0/0.32767          up    up   aenet    --> ae0.32767
ge-2/0/1.0              up    up   aenet    --> ae0.0
ge-2/0/1.32767          up    up   aenet    --> ae0.32767
ae0                     up    up
ae0.0                   up    up   bridge
ae0.32767               up    up   multiservice
```

The output shows the bundle relationship for the aggregated Ethernet interface and the overall status of the interface, including the following information:

- The link aggregation control PDUs run on the .0 child logical interfaces for the untagged aggregated Ethernet interface.
- The link aggregation control PDUs run on the .32767 child logical interfaces for the VLAN-tagged aggregated Ethernet interface.
- The .32767 logical interface is created for the parent link and all child links.

Verifying Aggregated Ethernet Interfaces (extensive)

IN THIS SECTION

- Purpose | 187
- Action | 187

Purpose

Display status information and statistics in extensive (detailed) format for aggregated Ethernet interfaces.

Action

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

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  ...
  Logical interface ae0.0 (Index 67) (SNMP ifIndex 628) (Generation 134)
    ...
    LACP info:        Role     System            System      Port    Port  Port
                             priority          identifier  priority  number   key
      ge-5/0/0.0     Actor        127  00:1f:12:8c:af:c0       127     832     1
      ge-5/0/0.0   Partner        127  00:1f:12:8f:d7:c0       127     640     1
      ge-5/0/1.0     Actor        127  00:1f:12:8c:af:c0       127     833     1
      ge-5/0/1.0   Partner        127  00:1f:12:8f:d7:c0       127     641     1
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-5/0/0.0             12830        7090            0            0
      ge-5/0/1.0             10304        4786            0            0
    ...
  Logical interface ae0.32767 (Index 70) (SNMP ifIndex 630) (Generation 135)
    ...
    LACP info:        Role     System            System      Port    Port  Port
                             priority          identifier  priority  number   key
      ge-5/0/0.32767    Actor      127 00:1f:12:8c:af:c0       127     832     1
      ge-5/0/0.32767  Partner      127 00:1f:12:8f:d7:c0       127     640     1
      ge-5/0/1.32767    Actor      127 00:1f:12:8c:af:c0       127     833     1
      ge-5/0/1.32767  Partner      127 00:1f:12:8f:d7:c0       127     641     1
```

```
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-5/0/0.32767         12830        7090            0            0
      ge-5/0/1.32767         10304        4786            0            0
  ...
```

The output shows detailed aggregated Ethernet interface information. This portion of the output shows LACP information and LACP statistics for each logical aggregated Ethernet interface.

RELATED DOCUMENTATION

Configuring Aggregated Ethernet Interfaces | 174

Release History Table

Release	Description
15.1X49-D40	Starting with Junos OS Release 15.1X49-D40, LACP is supported on Layer 2 transparent mode in addition to existing support on Layer 3 mode.

Configuring Link Aggregation Control Protocol

IN THIS SECTION

- Understanding LACP on Standalone Devices | 189
- Example: Configuring Link Aggregation Control Protocol | 189
- Verifying LACP on Standalone Devices | 195
- LAG and LACP Support Line Devices with I/O Cards (IOCs) | 198
- Example: Configuring LAG Interface on an Line Device with IOC2 or IOC3 | 200

Link Aggregation Control Protocol (LACP) provides a standard means for information exchange between the systems on a link. The below topics discuss the overview of LACP on standalone devices, examples of configuring LACP, LAG and LACP support line devices.

Understanding LACP on Standalone Devices

Link Aggregation Control Protocol (LACP) provides a standardized means for exchanging information between partner systems on a link. Within LACP, the local end of a child link is known as the actor and the remote end of the link is known as the partner.

LACP is enabled on an aggregated Ethernet interface by setting the mode to either passive or active. However, to initiate the transmission of link aggregation control protocol data units (PDUs) and response link aggregation control PDUs, you must enable LACP at both the local and remote ends of the links, and one end must be active:

- Active mode—If either the actor or partner is active, they exchange link aggregation control PDUs. The actor sends link aggregation control PDUs to its protocol partner that convey what the actor knows about its own state and that of the partner's state.
- Passive mode—If the actor and partner are both in passive mode, they do not exchange link aggregation control PDUs. As a result, the aggregated Ethernet links do not come up. In passive transmission mode, links send out link aggregation control PDUs only when they receive them from the remote end of the same link.

By default, the actor and partner transmit link aggregation control PDUs every second. 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.

You configure the interval at which the interfaces on the remote side of the link transmit link aggregation control PDUs by configuring the `periodic` statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval. The interval can be `fast` (every second) or `slow` (every 30 seconds).

NOTE: Starting with Junos OS Release 15.1X49-D40, LACP is supported in Layer 2 transparent mode in addition to existing support in Layer 3 mode.

Example: Configuring Link Aggregation Control Protocol

IN THIS SECTION

Requirements | **190**

- Overview | 190
- Configuration | 190
- Verification | 193

This example shows how to configure LACP.

Requirements

This example uses an SRX Series device.

Before you begin:

- Determine which interfaces to use and verify that they are in switch mode. See Understanding VLANs.

Overview

In this example, for aggregated Ethernet interfaces, you configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface.

Configuration

IN THIS SECTION

- Procedure | 190

Procedure

CLI Quick Configuration

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

```
set interfaces ge-0/0/6 ether-options 802.3ad ae0
set interfaces ge-0/0/7 ether-options 802.3ad ae0
```

```
set interfaces ae0 vlan-tagging
set interfaces ae0 aggregated-ether-options lacp active periodic fast
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set vlan vlan1000 vlan-id 1000
set interfaces ae0 unit 0 family ethernet-switching vlan members vlan1000
```

Step-by-Step Procedure

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

To configure LACP:

1. Configure the interfaces for ae0.

```
[edit ]
user@host# set interfaces ge-0/0/6 ether-options 802.3ad ae0
user@host# set interfaces ge-0/0/7 ether-options 802.3ad ae0
```

2. Configure ae0 interface for vlan tagging.

```
[edit ]
user@host# set interfaces ae0 vlan-tagging
```

3. Configure LACP for ae0 and configure periodic transmission of LACP packets.

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

4. Configure ae0 as a trunk port.

```
[edit ]
user@host# set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
```

5. Configure the VLAN.

```
[edit ]
user@host# set vlan vlan1000 vlan-id 1000
```

6. Add the ae0 interface to the VLAN.

```
[edit ]
user@host# set interfaces ae0 unit 0 family ethernet-switching vlan members vlan1000
```

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

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

Results

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

```
[edit]
user@host# show interfaces
ge-0/0/6 {
    ether-options {
        802.3ad ae0;
    }
}
ge-0/0/7 {
    ether-options {
        802.3ad ae0;
    }
}
ae0 {
    vlan- tagging;
    aggregated-ether-options {
        lacp {
            active;
            periodic fast;
```

```
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members vlan1000;
            }
        }
    }
}
```

Verification

IN THIS SECTION

- Verifying LACP Statistics | **193**
- Verifying LACP Aggregated Ethernet Interfaces | **194**

Verifying LACP Statistics

Purpose

Display LACP statistics for aggregated Ethernet interfaces.

Action

From operational mode, enter the `show lacp statistics interfaces ae0` command.

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-0/0/6                1352        2035            0           0
      ge-0/0/7                1352        2056            0           0
```

Meaning

The output shows LACP statistics for each physical interface associated with the aggregated Ethernet interface, such as the following:

- The LACP received counter that increments for each normal hello packet received
- The number of LACP transmit packet errors logged
- The number of unrecognized packet errors logged
- The number of invalid packets received

Use the following command to clear the statistics and see only new changes:

```
user@host# clear lacp statistics interfaces ae0
```

Verifying LACP Aggregated Ethernet Interfaces

Purpose

Display LACP status information for aggregated Ethernet interfaces.

Action

From operational mode, enter the `show lacp interfaces ae0` command.

```
user@host> show lacp interfaces ae0
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/0/6       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/6     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Passive
      ge-0/0/7       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/7     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Passive
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-0/0/6                  Current   Fast periodic Collecting distributing
      ge-0/0/7                  Current   Fast periodic Collecting distributing
```

Meaning

The output shows aggregated Ethernet interface information, including the following information:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up if it is collecting and distributing packets.

SEE ALSO

Understanding Link Aggregation Control Protocol

Ethernet Ports Switching Overview for Security Devices

Verifying LACP on Standalone Devices

IN THIS SECTION

- Verifying LACP Statistics | **195**
- Verifying LACP Aggregated Ethernet Interfaces | **196**

Verifying LACP Statistics

IN THIS SECTION

- Purpose | **195**
- Action | **196**

Purpose

Display LACP statistics for aggregated Ethernet interfaces.

Action

From operational mode, enter the `show lacp statistics interfaces ae0` command.

```
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
```

The output shows LACP statistics for each physical interface associated with the aggregated Ethernet interface, such as the following:

- The LACP received counter that increments for each normal hello
- The number of LACP transmit packet errors logged
- The number of unrecognized packet errors logged
- The number of invalid packets received

Use the following command to clear the statistics and see only new changes:

```
user@host# clear lacp statistics interfaces ae0
```

Verifying LACP Aggregated Ethernet Interfaces

IN THIS SECTION

- Purpose | **196**
- Action | **197**

Purpose

Display LACP status information for aggregated Ethernet interfaces.

Action

From operational mode, enter the `show lacp interfaces ae0` command.

```
user@host> show lacp interfaces ae0
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-2/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-2/0/0                  Current   Fast periodic Collecting distributing
      ge-2/0/1                  Current   Fast periodic Collecting distributing
      ge-2/2/0                  Current   Fast periodic Collecting distributing
      ge-2/2/1                  Current   Fast periodic Collecting distributing
```

The output shows aggregated Ethernet interface information, including the following information:

- The LACP state—Indicates whether the link in the bundle is an actor (local or near-end of the link) or a partner (remote or far-end of the link).
- The LACP mode—Indicates whether both ends of the aggregated Ethernet interface are enabled (active or passive)—at least one end of the bundle must be active.
- The periodic link aggregation control PDU transmit rate.
- The LACP protocol state—Indicates the link is up ifit is collecting and distributing packets.

RELATED DOCUMENTATION

Verifying LACP on Redundant Ethernet Interfaces

LAG and LACP Support Line Devices with I/O Cards (IOCs)

IN THIS SECTION

- LAG and LACP Support on the SRX5000 Module Port Concentrator | 198
- LAG and LACP Support on the SRX5000 Line IOCs in Express Path Mode | 199

NOTE: The following notes apply to 'LAG and LACP Support on SRX5000 Line Devices' as outlined in this document.

- Cross-IOC LAG interfaces do not support Layer 2 transparent mode.
- Mixed interface speeds are supported on the same aggregated bundle.
- A redundant Ethernet interface or aggregated Ethernet interface must contain child interfaces from the same IOC type. For example, if one child link is from 10-Gigabit Ethernet on IOC2, the second child link should also be from IOC2. Similarly, both child interfaces can be from IOC3. Configuring child interfaces by mixing links from both IOC2 and IOC3 is not supported.

LAG and LACP Support on the SRX5000 Module Port Concentrator

The SRX5000 Module Port Concentrator (SRX5K-MPC) on SRX5400, SRX5600, and SRX5800 devices supports link aggregation groups (LAGs) and Link Aggregation Control Protocol (LACP).

Support for LAGs based on IEEE 802.3ad makes it possible to aggregate physical interface links on your device. LAGs provide increased interface bandwidth and link availability by linking physical ports and load-balancing traffic crossing the combined interface.

LACP provides a standardized means for exchanging information between partner (remote or far-end of the link) systems on a link. This exchange allows their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG. This exchange also enables the transmission and reception processes for the link to function in an orderly manner.

The following LAG and LACP features are supported on the SRX5K-MPC:

- Bandwidth aggregation—Increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

- Link redundancy and load balancing (within chassis cluster)—Provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.
- Dynamic link management—Enables automatic addition and deletion of individual links to the aggregate bundle without user intervention.

LACP supports the following features:

- LACP bundles several physical interfaces to form one logical interface by exchanging LACP packets between the local interface and the remote interface. LACP monitors the link for changes in interface state by exchanging a periodic LACP heartbeat between two sides. Any changes in interface state are reflected in the LACP packet.
- Normally after an LACP is configured and committed, two sides start to exchange interface and port information. Once they identify each other and match the LACP state machine criteria, the LACP is declared as up. You can deactivate or delete the LACP configuration.
- By default, the LACP packets are exchanged in every second. You can configure the LACP interval as fast (every second) or slow (every 30 seconds) to ensure the health of the interfaces.
- LACP supports distributed and centralized modes. Chassis cluster setup is recommended to operate with LACP distributed mode, which handles chassis cluster failover better. The centralized mode might experience traffic loss during failover.

SRX5K-MPCs on SRX5000 line devices provide active and standby support with redundant Ethernet interface LAGs in chassis cluster deployments.

LAG and LACP Support on the SRX5000 Line IOCs in Express Path Mode

Starting in Junos OS Release 15.1X49-D40, the IOC2 and IOC3 cards on SRX5400, SRX5600, and SRX5800 devices support link aggregation groups (LAGs) and Link Aggregation Control Protocol (LACP) in Express Path mode.

You can use the links in a LAG as ingress or egress interfaces in Express Path mode. The LAG links can include links from cards such as IOC2 or IOC3. For a LAG link to qualify for Express Path, all its member links should be connected to Express Path-enabled network processors. If Express Path is disabled on any of the member links in a LAG, a regular session (non-Express Path session) is created.

NOTE:

- Cross-IOC LAG interfaces do not support Layer 2 transparent mode.

- Mixed interface speeds are supported on the same aggregated bundle.
- A redundant Ethernet interface or aggregated Ethernet interface must contain child interfaces from the same IOC type. For example, if one child link is from 10-Gigabit Ethernet on IOC2, the second child link should also be from IOC2. Similarly, both child interfaces can be from IOC3. Configuring child interfaces by mixing links from both IOC2 and IOC3 is not supported.

SEE ALSO

Configuring Aggregated Ethernet Interfaces | **174**

Configuring Link Aggregation Control Protocol | **188**

Example: Configuring LACP on Chassis Clusters

Example: Configuring LAG Interface on an Line Device with IOC2 or IOC3

IN THIS SECTION

- Requirements | **200**
- Overview | **201**
- Configuration | **201**
- Verification | **205**

Starting in Junos OS Release 15.15X49-D40, IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface. This single, aggregated Ethernet interface is also known as a LAG or bundle. The LACP provides additional functionality for LAGs.

This example shows how to configure LAG on an SRX Series device using the links from either IOC2 or IOC3 in Express Path mode.

Requirements

This example uses the following software and hardware components:

- Junos OS Release 15.1X49-D40 or later for SRX Series devices.

- SRX5800 with IOC2 or IOC3 with Express Path enabled on IOC2 and IOC3. For details, see *Express Path*.

Overview

In this example, you create a logical aggregated Ethernet interface and define the parameters associated with the logical aggregated Ethernet interface, such as a logical unit, interface properties, and LACP. Next, define the member links to be contained within the aggregated Ethernet interface—for example, four 10-Gigabit Ethernet interfaces. Finally, configure an LACP for link detection.

The following member links are used in this example:

- xe-0/0/8
- xe-0/0/9
- xe-1/0/8
- xe-1/0/9
- xe-3/1/4
- xe-3/1/5
- xe-5/1/4
- xe-5/1/5

Configuration

IN THIS SECTION

- CLI Quick Configuration | **201**
- Procedure | **202**

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, delete, and then copy

and paste the commands into the CLI at the [edit] hierarchy level, and then enter commit from configuration mode.

```
set chassis aggregated-devices ethernet device-count 5
set interfaces xe-0/0/8 gigether-options 802.3ad ae1
set interfaces xe-0/0/9 gigether-options 802.3ad ae0
set interfaces xe-1/0/8 gigether-options 802.3ad ae1
set interfaces xe-1/0/9 gigether-options 802.3ad ae0
set interfaces xe-3/1/4 gigether-options 802.3ad ae1
set interfaces xe-3/1/5 gigether-options 802.3ad ae0
set interfaces xe-5/1/4 gigether-options 802.3ad ae1
set interfaces xe-5/1/5 gigether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 17.0.0.1/24
set interfaces ae1 unit 0 family inet address 16.0.0.1/24
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp active
```

Procedure

Step-by-Step Procedure

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

To configure LAG Interfaces:

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

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

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

```
[edit interfaces]
user@host# set xe-0/0/8 gigether-options 802.3ad ae1
user@host# set xe-0/0/9 gigether-options 802.3ad ae0
user@host# set xe-1/0/8 gigether-options 802.3ad ae1
user@host# set xe-1/0/9 gigether-options 802.3ad ae0
user@host# set xe-3/1/4 gigether-options 802.3ad ae1
user@host# set xe-3/1/5 gigether-options 802.3ad ae0
```

```
user@host# set xe-5/1/4 gigether-options 802.3ad ae1
user@host# set xe-5/1/5 gigether-options 802.3ad ae0
```

3. Assign an IP address to ae0 and ae1.

```
[edit interfaces]
user@host# set ae0 unit 0 family inet address 17.0.0.1/24
user@host# set ae1 unit 0 family inet address 16.0.0.1/24
```

4. Set the LACP on reth0.

```
[edit interfaces]
user@host# set ae0 aggregated-ether-options lacp active
user@host# set ae1 aggregated-ether-options lacp active
```

Results

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

```
[edit]
user@host# show interfaces
xe-0/0/8 {
    gigether-options {
        802.3ad ae1;
    }
}
xe-0/0/9 {
    gigether-options {
        802.3ad ae0;
    }
}
xe-1/0/8 {
    gigether-options {
        802.3ad ae1;
    }
}
xe-1/0/9 {
```

```
    gigether-options {
        802.3ad ae0;
    }
}
xe-3/1/4 {
    gigether-options {
        802.3ad ae1;
    }
}
xe-3/1/5 {
    gigether-options {
        802.3ad ae0;
    }
}
ae0 {
    aggregated-ether-options {
        lacp {
            active;
        }
    }
    unit 0 {
        family inet {
            address 17.0.0.1/24;
        }
    }
}
ae1 {
    aggregated-ether-options {
        lacp {
            active;
        }
    }
    unit 0 {
        family inet {
            address 16.0.0.1/24;
        }
    }
}
```

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

```
aggregated-devices {
    ethernet {
        device-count 5;
    }
}
```

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

Verification

IN THIS SECTION

- Verifying LACP on Redundant Ethernet Interfaces | **205**

Verifying LACP on Redundant Ethernet Interfaces

Purpose

Display LACP status information for redundant Ethernet interfaces.

Action

From operational mode, enter the `show lacp interfaces` command to check that LACP has been enabled as active on one end.

```
user@host> show lacp interfaces
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/9       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/9     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/9       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/9     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/5       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/5     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/5       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/5     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/9                  Current   Fast periodic Collecting distributing
      xe-1/0/9                  Current   Fast periodic Collecting distributing
```

```
      xe-3/1/5                Current   Fast periodic Collecting distributing
      xe-5/1/5                Current   Fast periodic Collecting distributing

Aggregated interface: ae1
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      xe-0/0/8       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-0/0/8     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/8       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-1/0/8     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-3/1/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/4       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      xe-5/1/4     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      xe-0/0/8                Current   Fast periodic Collecting distributing
      xe-1/0/8                Current   Fast periodic Collecting distributing
      xe-3/1/4                Current   Fast periodic Collecting distributing
      xe-5/1/4                Current   Fast periodic Collecting distributing
```

The output indicates that LACP has been set up correctly and is active at one end.

Release History Table

Release	Description
15.1X49-D40	Starting in Junos OS Release 15.1X49-D40, the IOC2 and IOC3 cards on SRX5400, SRX5600, and SRX5800 devices support link aggregation groups (LAGs) and Link Aggregation Control Protocol (LACP) in Express Path mode.
15.1X49-D40	Starting in Junos OS Release 15.15X49-D40, IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface.

Configuring Gigabit Ethernet Physical Interface Modules

IN THIS SECTION

- Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM | 207
- Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface | 210
- Understanding the 2-Port 10-Gigabit Ethernet XPIM | 219
- Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface | 222
- Understanding the 8-Port Gigabit Ethernet SFP XPIM | 228
- Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs | 230

Small form-factor pluggables (SFPs) are hot-pluggable modular interface transceivers for Gigabit and Fast Ethernet connections. The 1-Port Gigabit Ethernet SFP Mini-PIM interfaces a single Gigabit Ethernet device or a network. The below topics discuss the overview and configuration of 1-Port Gigabit Ethernet SFP Mini-PIM interface, overview and configuration of 2-Port 10-GE XPIM and overview and configuration of 8-Port GE SFP XPIMs.

Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM

IN THIS SECTION

- Supported Features | 208
- Interface Names and Settings | 208
- Available Link Speeds and Modes | 209
- Link Settings | 209

Small form-factor pluggables (SFPs) are hot-pluggable modular interface transceivers for Gigabit and Fast Ethernet connections. Gigabit Ethernet SFP Mini-PIMs can be used in copper and optical

environments to provide maximum flexibility when upgrading from an existing infrastructure to Metro Ethernet.

The 1-Port Gigabit Ethernet SFP Mini-PIM interfaces a single Gigabit Ethernet device or a network. It supports a variety of transceivers with data speeds of 10-Mbps/100-Mbps/1-Gbps with extended LAN or WAN connectivity.

Transceivers are hot-swappable.

This topic includes the following sections:

Supported Features

The following features are supported on the 1-Port Gigabit Ethernet SFP Mini-PIM:

- 10-Mbps/100-Mbps/1-Gbps link speed
- Half-duplex/full-duplex support
- Autonegotiation
- Encapsulations
- Maximum transmission unit (MTU) size of 1514 bytes (default) and 9010 bytes (jumbo frames)
- Loopback
- Transceivers are hot-swappable

Interface Names and Settings

The following format is used to represent the 1-Port Gigabit Ethernet SFP Mini-PIM interface names:

type-fpc/pic/port

Where:

- type—Media type (ge)
- fpc—Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- pic—Number of the PIC on which the physical interface is located (0)
- port—Specific port on a PIC (0)

Examples: `ge-1/0/0` and `ge-2/0/0`

By default, the interfaces on the ports on the uplink module installed on the device are enabled. You can also specify the MTU size for the Gigabit Ethernet interface. Junos OS supports values from 256 through 9010. The default MTU size for Gigabit Ethernet interfaces is 1514.

Available Link Speeds and Modes

The 1-Port Gigabit Ethernet SFP Mini-PIM supports the following link speeds:

- `10m`—Sets the link speed to 10 Mbps.
- `100m`—Sets the link speed to 100 Mbps.
- `1g`—Sets the link speed to 1 Gbps.

The 1-Port Gigabit Ethernet SFP Mini-PIM supports the following link modes:

- `Full-duplex`—Allows bidirectional communication at a given point in time.
- `Half-duplex`—Allows single directional communication at a given point in time.

Link Settings

The 1-Port Gigabit Ethernet SFP Mini-PIM includes the following link settings:

- `auto-negotiation`—Enables autonegotiation of link mode and speed.

  NOTE: By default, autonegotiation is enabled. To disable autonegotiation, use `set gigether-options no-autonegotiation`

  We recommend enabling autonegotiation.

- `loopback`—Enables loopback.
- `no-auto-negotiation`—Disables autonegotiation of link mode and speed.
- `no-loopback`—Disables loopback.

By default a link speed of 1 Gbps in full-duplex mode is supported.

NOTE: On SRX340 High Memory devices, traffic might stop between the SRX340 device and the Cisco switch due to link mode mismatch. We recommend setting the same value to the autonegotiation parameters on both ends.

NOTE: On SRX300 devices, the link goes down when you upgrade FPGA on 1-Port Gigabit Ethernet SFP mini-PIM. As a workaround, run the `restart fpc` command and restart the FPC.

Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface

IN THIS SECTION

- Requirements | **210**
- Overview | **210**
- Configuration | **211**
- Verification | **215**

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

Requirements

Before you begin:

- Establish basic connectivity. See the *Getting Started Guide* for your device.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 166.

Overview

In this example, you configure the ge-2/0/0 interface, set the operating speed to 100 Mbps, and define a logical interface that you can connect to the 1-Port Gigabit Ethernet SFP Mini-PIM. You also set the MTU value to 9010 and set the link option to no-loopback.

Configuration

IN THIS SECTION

- Procedure | 211
- Configuring Physical Properties | 211
- Disabling the Interface | 212
- Configuring Logical Properties | 212
- Editing Logical Properties | 213
- Deleting the Logical Interface | 213
- Configuring a 1-Port Gigabit Ethernet SFP Mini-PIM | 214

Procedure

CLI Quick Configuration

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

```
set interfaces ge-2/0/0 link-mode full-duplex speed 100m
set interface ge-2/0/0 gigether-options no-loopback
```

Configuring Physical Properties

GUI Quick Configuration

Step-by-Step Procedure

To quickly configure the physical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select `Configure > Interfaces`.
2. Under Interface, select `ge-2/0/0` and then click `Edit`. A pop-up window appears.

3. In the Description box, type the description for the SFP Mini-PIM.
4. In the MTU box, type `9010`.
5. From the Speed list, select `100Mbps`.
6. From the Link-mode list, select `Full-duplex`.
7. Select the Enable Auto-negotiation checkbox.
8. Select the Enable Per Unit Scheduler checkbox.
9. Click `OK`

Disabling the Interface

GUI Quick Configuration

Step-by-Step Procedure

To disable the 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select `Configure > Interfaces` .
2. Under Interface, select `ge-2/0/0` and then click `Disable`.

Configuring Logical Properties

GUI Quick Configuration

Step-by-Step Procedure

To quickly configure the logical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web, use the following steps:

1. Select `Configure > Interfaces`.
2. Under Interface, select `ge-2/0/0.0`, and then click `Add Logical Interface`. A pop-up window appears.
3. In the Unit box, type `0`.
4. In the Description box, type a description for the SFP Mini-PIM.
5. From the Zone list, select `untrust`.

6. To edit the family protocol type to the Mini-PIM interfaces, select the IPv4 tab, and then select `Enable address configuration`.
7. Click `Add`, and then type IPv4 address.
8. Click `OK`.

Editing Logical Properties

Step-by-Step Procedure

To quickly configure the physical properties of a 1-Port Gigabit Ethernet SFP Mini-PIM using J-Web:

1. Under Interface, select the logical interface added to the 1-Port Gigabit Ethernet SFP Mini-PIM and then click `Edit`. A pop-up window appears.
2. Under Interface, select `ge-2/0/0.0`, and then click `Edit Logical Interface`. A pop-up window appears.
3. From the Zone list, select `trust`.
4. To enable DHCP client on the interface, select the IPv4 tab and then select `Enable DHCP`.
5. Click `OK`.

> **NOTE**: You cannot add or edit Description and Unit for a logical interface.

Deleting the Logical Interface

GUI Quick Configuration

Step-by-Step Procedure

To delete the logical interface of 1–Port Gigabit Ethernet SFP Mini-PIM using J-Web,

1. Select `Configure > Interfaces`.
2. Under Interface, select `ge-2/0/0.0`, and then click `Delete`.

Configuring a 1-Port Gigabit Ethernet SFP Mini-PIM

Step-by-Step Procedure

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

To configure a 1-Port Gigabit Ethernet SFP Mini-PIM:

1. Configure the interface.

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

2. Set the operating link-mode full-duplex speed of 100 Mbps for the SFP Mini-PIM.

```
[edit interfaces ge-2/0/0]
user@host# set link-mode full-duplex speed 100m
```

3. Assign the MTU value.

```
[edit interfaces ge-2/0/0]
user@host# set mtu 9010
```

4. Add the logical interface.

```
[edit interfaces ge-2/0/0]
user@host# set unit 0 family inet address 14.1.1.1/24
```

5. Set the link options.

```
[edit interfaces ge-2/0/0]
user@host# set gigether-options no-loopback
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces ge-2/0/0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces ge-2/0/0
    mtu 9010;
    speed 100m;
    gigether-options {
    no-loopback;
}
unit 0 {
    family inet {
    14.1.1.1/24
    }
}
```

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

Verification

IN THIS SECTION

- Verifying That the Correct Hardware Is Installed | **215**
- Verifying the FPC Status | **217**
- Verifying the Interface Settings | **217**

Confirm that the configuration is working properly.

Verifying That the Correct Hardware Is Installed

Purpose

Verify that the 1-Port Gigabit Ethernet SFP Mini-PIM is installed on the device.

Action

From operational mode, enter the `show chassis hardware` command.

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                AG0309AA0004      SRX240b
Routing Engine   REV 16   750-021792   VL3180            RE-SRX240B
  da0     999 MB  ST72682                                Nand Flash
  usb0 (addr 1)  DWC OTG root hub 0    vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x005a 90  vendor 0x0409     uhub1
  usb0 (addr 3)  ST72682  High Speed Mode 64218 STMicroelectronics umass0
FPC 0                                                  FPC
  PIC 0                                                16x GE Base PIC
FPC 1                    750-023367   112009000278      FPC
  PIC 0                                                1x T1E1 mPIM
FPC 2            REV 00   750-03273   AABC5081          FPC
  PIC 0                                                1x GE High-Perf SFP mPIM
    Xcvr 0       REV 02   740-011612   9101465           SFP-T
FPC 4                    750-029145   122009000061      FPC
PIC 0                                                    1x GE SFP mPIM
    Xcvr 0       REV 01   740-011782   PBL0C3T           SFP-SX
Power Supply 0
```

Verify that the output contains the following values:

- `FPC 2, PIC 0` —1x GE High-Perf SFP mPIM
- `FPC 4, PIC 0` —1x GE SFP mPIM

NOTE: In the example shown above, the output for 1-Port SFP Mini-Physical Interface Module is displayed as 1X GE SFP mPIM and the output for 1-Port Gigabit Ethernet SFP Mini-Physical Interface Module is displayed as 1X GE High-Perf SFP mPIM.

NOTE: The 1-Port GE SFP Mini-PIM is installed in the second slot of the device chassis; therefore the output displayed is 1x GE High-Perf SFP mPIM and the Flexible PIC Concentrator (FPC) used here is fpc 2.

The 1-Port SFP Mini-PIM is installed in the fourth slot of the device chassis; therefore the output displayed is 1x GE SFP mPIM and Flexible PIC Concentrator (FPC) used here is fpc 4.

Verifying the FPC Status

Purpose

Verify the FPC status.

Action

From operational mode, enter the `show chassis fpc` command.

```
show@host> show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State           (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online          -------------------- CPU less FPC --------------------
  1  Online          -------------------- CPU less FPC --------------------
  2  Online          -------------------- CPU less FPC --------------------
  3  Empty
  4  Online          -------------------- CPU less FPC --------------------
```

The output should show the FPC status as online.

The 1-Port SFP Mini-PIM is installed in the fourth slot of the device chassis; the output shows the FPC status for slot 4 as online.

The 1-Port Gigabit Ethernet SFP Mini-PIM is installed in the second slot of the device chassis; the output shows the FPC status for slot 2 as online.

Verifying the Interface Settings

Purpose

Verify that the interface is configured as expected.

Action

From operational mode, enter the `show interface ge-2/0/0` command.

```
user@host# run show interfaces ge-2/0/0
Physical interface: ge-2/0/0, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 552
  Link-level type: Ethernet, MTU: 9010, Link-mode: Full-duplex, Speed: 100mbps, 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
  Current address: 00:22:83:99:ac:f2, Hardware address: 00:22:83:99:ac:f2
  Last flapped   : 2010-08-17 12:20:33 UTC (00:00:20 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-2/0/0.0 (Index 88) (SNMP ifIndex 557)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 108
    Output packets: 1
    Security: Zone: Null
    Protocol inet, MTU: 8996
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 14.1.1/24, Local: 14.1.1.1, Broadcast: 14.1.1.255
```

Verify the following information in the command output:

- Physical interface—ge-2/0/0, Enabled, Physical link is Up
- MTU—9010; Link-mode—Full-duplex
- Speed—100 Mbps
- Loopback—Disabled

Understanding the 2-Port 10-Gigabit Ethernet XPIM

IN THIS SECTION

- Supported Features | 219
- Interface Names and Settings | 220
- Copper and Fiber Operating Modes | 221
- Link Speeds | 221
- Link Settings | 221

The 10-Gigabit Ethernet (also known as 10GBASE-T or IEEE 802.3an) is a telecommunication technology that offers data speeds up to 10 billion bits per second over unshielded or shielded twisted pair cables.

The 2-Port 10-Gigabit Ethernet *Physical Interface Module* (XPIM) is a 2 x 10GBASE-T / SFP+ XPIM line card. (SFP+ is a fiber optic transceiver module designed for 10-Gigabit Ethernet and 8.5 Gbps-fiber channel systems.) The 2-Port 10-Gigabit Ethernet XPIM provides a front-end interface connection that includes the following ports:

- 2 X copper ports. The copper ports support 10GBASE-T running with CAT6A or CAT7 Ethernet cable for up to 100 meters.
- 2 X fiber (SFP+) ports. The fiber ports support SFP+ multiple 10G modules.

The 2-Port 10-Gigabit Ethernet XPIM provides interconnects for LANs, WANs, and metropolitan area networks (MANs). The XPIM provides multiple service levels (1-Gigabit Ethernet to 10-Gigabit Ethernet in increments) and a single connection option for a wide range of customer needs and applications.

NOTE: By default, the 2-Port 10-Gigabit Ethernet XPIM ports comes up in fiber mode, while autonegotiation is not supported.

This topic includes the following sections:

Supported Features

The following features are supported on the 2-Port 10-Gigabit Ethernet XPIM:

- Multiple SFP+ 10G modules and the following SFP modules:

  - SFPP-10GE-SR
  - SFPP-10GE-LR
  - SFPP-10GE-ER
  - SFPP-10GE-LRM

- Copper TWIN-AX 1M and Copper TWIN-AX 3M
- Online Insertion and Removal (OIR) functionality
- Link speeds of up to 10-Gbps
- Full-duplex and half-duplex modes
- Flow control
- Autonegotiation and autosensing
- *Quality of service* (QoS)

Interface Names and Settings

The following format is used to represent the 2-Port 10-Gigabit Ethernet XPIM interface names:

type-fpc/pic/port

Where:

- type — Media type (xe)
- fpc — Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- pic — Number of the PIC on which the physical interface is located (0)
- port — Specific port on a PIC (0 or 1)

By default, the interfaces (for example, `xe-6/0/0` or `xe-2/0/0`) on the ports on the uplink module installed on the device are enabled. You can also specify the maximum transmission unit (MTU) size for the Gigabit Ethernet interface. Junos OS supports values from 256 through 9192. The default MTU for Gigabit Ethernet interfaces is 1514.

Copper and Fiber Operating Modes

On the 2-Port 10-Gigabit Ethernet XPIM, one copper port and one fiber port is grouped together as port 0, and another copper port and fiber port are grouped as port 1. Only two ports can be active at the same time (one port from port 0 and another port from port 1).

The 2-Port 10-Gigabit Ethernet XPIM can be configured to operate in two copper mode, two fiber mode, or mixed mode (one copper and one fiber). In mixed mode, the two ports should be from different port groups (one port from port 1 and the other from port 2).

Link Speeds

The 2-Port 10-Gigabit Ethernet XPIM ports support the following link speeds for copper and fiber:

- Copper—10/100/1000 Mbps or 10Gbps (full duplex). Half-duplex is only for 10/100 Mbps.
- Fiber—1000 Mbps or 10 Gbps (full duplex). Half-duplex mode is not supported.

To set the link speeds, use the following options:

- `10m`—Sets the link speed to 10 Mbps.
- `10g`—Sets the link speed to 10 Gbps.
- `100m`—Sets the link speed to 100 Mbps.
- `1g`—Sets the link speed to 1 Gbps.

Link Settings

The 2-Port 10-Gigabit Ethernet XPIM includes the following link settings:

- `802.3ad`—Specifies an aggregated Ethernet bundle.
- `auto-negotiation`—Enables autonegotiation of flow control, link mode, and speed.
- `loopback`—Enables loopback.
- `no-auto-negotiation`—Disables autonegotiation of flow control, link mode, and speed.
- `no-loopback`—Disables loopback.

By default, flow control is enabled on all ports, a link speed of 10 Gbps in full duplex is supported, autonegotiation is disabled on the fiber ports, and autonegotiation is enabled on copper ports.

NOTE: Autonegotiation is not supported when the 2-Port 10-Gigabit Ethernet XPIM is operating in fiber mode at a link speed of 10 Gbps.

Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface

IN THIS SECTION

- Requirements | 222
- Overview | 222
- Configuration | 223
- Verification | 225

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

Requirements

Before you begin:

- Establish basic connectivity. See the *Getting Started Guide* for your device.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 166.

Overview

In this example, you configure the xe-6/0/0 interface, set the operating mode to copper mode, set the operating speed to 10 Gbps, and define a logical interface that you can connect to the 2-Port 10-Gigabit Ethernet XPIM. Additionally, you set the MTU value to 1514, set the link option to no loopback, and enable the interface.

Configuration

IN THIS SECTION

- Procedure | 223

Procedure

CLI Quick Configuration

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

```
set interfaces xe-6/0/0 media-type copper speed 10g unit 0 family inet mtu 1514
set interface xe-6/0/0 gigether-options no-loopback
```

Step-by-Step Procedure

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

To configure a 2-Port 10-Gigabit Ethernet XPIM:

1. Configure the interface.

```
[edit]
user@host# edit interfaces xe-6/0/0
```

2. Configure the operating mode.

```
[edit interfaces xe-6/0/0]
user@host# set media-type copper
```

3. Set the operating speed for the XPIM.

```
[edit interfaces xe-6/0/0]
user@host# set speed 10g
```

4. Add the logical interface.

```
[edit interfaces xe-6/0/0]
user@host# set unit 0 family inet
```

5. Assign the physical interface MTU value.

```
[edit interfaces xe-6/0/0]
user@host# set interface xe-6/0/0 mtu 1514
```

6. Assign the logical interface MTU value.

```
[edit interfaces xe-6/0/0]
user@host# set unit 0 family inet mtu 1500
```

7. Set the link options.

```
[edit interfaces xe-6/0/0]
user@host# set gigether-options no-loopback
```

8. Disable the interface.

```
[edit interfaces xe-6/0/0]
user@host# set disable
```

9. Enable the interface.

```
[edit interfaces xe-6/0/0]
user@host# delete disable
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces xe-6/0/0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces xe-6/0/0
    speed 10g;
    media-type copper;
    gigether-options {
    no-loopback;
}
unit 0 {
    family inet {
    mtu 1514;
    }
}
```

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

Verification

IN THIS SECTION

- Verifying That the Correct Hardware Is Installed | 225
- Verifying the FPC Status | 226
- Verifying the Interface Settings | 227

Confirm that the configuration is working properly.

Verifying That the Correct Hardware Is Installed

Purpose

Verify that the 2-Port 10-Gigabit Ethernet XPIM is installed on the device.

Action

From operational mode, enter the `show chassis hardware` command.

```
Hardware inventory:
Item                      Version         Part number     Serial number         Description
Chassis                                                     AJ0309AC0047         SRX650
Midplane             REV 04         710-023875         TV3993
System IO             REV 04         710-023209         TV4035                SRXSME System IO
Routing Engine     REV 01         710-023224         DT5109                RE-SRXSME-SRE6
FPC 0                                                                            FPC
PIC 0                                                                            4x GE Base PIC
FPC 2                                                                            FPC
PIC 0                                                                            2x 10G gPIM
FPC 6                                                                            FPC
PIC 0                                                                            2x 10G gPIM
Power Supply 0     REV 01         740-024283         TA00049WSSSS         PS 645W AC
```

Verify that the output contains the following values:

- `FPC 2 , PIC 0`—2x 10G gPIM
- `FPC 6, PIC 0`—2x 10G gPIM

Verifying the FPC Status

Purpose

Verify the FPC status.

Action

From operational mode, enter the `show chassis fpc` command.

```
            Temp         CPU Utilization     (%)     Memory         Utilization (%)
Slot State         (C)         Total Interrupt             DRAM (MB)     Heap Buffer
0 Online         -------------------- CPU less FPC --------------------
1 Empty
2 Online         -------------------- CPU less FPC --------------------
3 Empty
4 Empty
```

```
5 Empty
6 Online         -------------------- CPU less FPC --------------------
7 Empty
8 Empty
```

The output should display FPC status as online.

Verifying the Interface Settings

Purpose

Verify that the interface is configured as expected.

Action

From operational mode, enter the `show interface xe-6/0/0` command.

```
Physical interface: xe-6/0/0, Enabled, Physical link is Up
    Interface index: 144, SNMP ifIndex: 501
    Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 10Gbps,
    BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
    Source filtering: Disabled, Flow control: Enabled
    Device flags : Present Running
    6 Copyright © 2010, Juniper Networks, Inc.
    Interface flags: SNMP-Traps Internal: 0x0
    Link flags : None
    CoS queues : 8 supported, 8 maximum usable queues
    Current address: 00:1f:12:e0:80:a8, Hardware address: 00:1f:12:e0:80:a8
    Last flapped : 1970-01-01 00:34:22 PST (07:26:29 ago)
    Input rate : 0 bps (0 pps)
    Output rate : 0 bps (0 pps)
    Active alarms : None
    Active defects : None

Logical interface xe-6/0/0.0 (Index 72) (SNMP ifIndex 503)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 25
    Output packets: 25
    Security: Zone: HOST
    Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
```

```
    ospf pgm pim rip router-discovery rsvp sap vrrp
    Protocol inet, MTU: 1500
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.10.10/24, Local: 10.10.10.10, Broadcast: 10.10.10.255
```

Verify the following information in the command output:

- Physical interface—xe-6/0/0, Enabled, Physical link is Up
- MTU—1514
- Link mode—Full duplex
- Speed—10 Gbps
- Loopback—Disabled
- Flow control—Enabled

Understanding the 8-Port Gigabit Ethernet SFP XPIM

IN THIS SECTION

- Supported Features | 229
- Interface Names and Settings | 230

A Gigabit Ethernet *Physical Interface Module* (XPIM) is a network interface card (NIC) that installs in the front slots of the SRX550 Services Gateway to provide physical connections to a LAN or a WAN.

NOTE: Starting in Junos OS Release 15.1X49-D10, the 8-Port Gigabit Ethernet SFP XPIM is not supported on legacy SRX Series systems. In Junos OS Release 15.1X49-D30, support for the 8-Port Gigabit Ethernet SFP XPIM is restored for SRX550 Service Gateway systems.

Small form-factor pluggables (SFPs) are hot-pluggable modular interface transceivers for gigabit and Fast Ethernet connections. The 8-port SFP Gigabit Ethernet interface enables customers to connect to Ethernet WAN services as well as to local servers at gigabit speed.

Supported Features

The following features are supported on the 8-Port Gigabit Ethernet SFP XPIM:

- Operates on both a slot with a maximum bandwidth of 8 gigabits and a slot with a maximum bandwidth of 1 gigabit
- Operates in tri-rate (10/100/1000 Mbps) mode with copper SFPs
- Routing and switched mode operation
- Layer 2 protocols
  - Link Aggregation Control Protocol (LACP)
  - Link Layer Discovery Protocol (LLDP)
  - GARP VLAN Registration Protocol (GVRP)
  - Internet Group Management Protocol (IGMP) snooping (v1 and v2)
  - Spanning Tree Protocol (STP), Real-Time Streaming Protocol (RTSP), and Multiple Spanning Tree Protocol (MSTP)
  - 802.1x
- Encapsulation (supported at the Physical Layer)
  - ethernet-bridge
  - ethernet-ccc
  - ethernet-tcc
  - ethernet-vpls
  - extended-vlan-ccc
  - extended-vlan-tcc
  - flexible-ethernet-services
  - vlan-ccc
- Q in Q VLAN tagging
- Integrated routing and bridging (IRB)
- Jumbo frames (9192 byte size)

- *Chassis cluster* switching
- Chassis cluster fabric link using GE ports

NOTE: The following Layer 2 switching features are not supported when the 8-Port Gigabit Ethernet SFP XPIM is plugged in slots with speeds of less than 1 gigabit:

- Q in Q VLAN tagging
- Link aggregation using ports across multiple XPIMs

Interface Names and Settings

The following format is used to represent the 8-Port SFP XPIM:

type-fpc/pic/port

Where:

- type—Media type (ge)
- fpc—Number of the Flexible PIC Concentrator (FPC) card where the physical interface resides
- pic—Number of the PIC where the physical interface resides (0)
- port—Specific port on a PIC (0)

Examples: `ge-1/0/0` and `ge-2/0/0`

By default, the interfaces on the ports on the uplink module installed on the device are enabled. You can also specify the maximum transmission unit (MTU) size for the XPIM. Junos OS supports values from 256 through 9192. The default MTU size for the 8-Port Gigabit Ethernet SFP XPIM is 1514.

Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs

IN THIS SECTION

- Requirements | 231
- Overview and Topology | 231
- Configuration | 232

- Verification | 238

This example shows how to perform a basic back-to-back device configuration with 8-port Gigabit Ethernet small form-factor pluggable (SFP) XPIMs. It describes a common scenario in which SFP XPIMs are deployed.

NOTE: Starting in Junos OS Release 15.1X49-D10, the 8-Port Gigabit Ethernet SFP XPIM is not supported on legacy SRX Series systems. In Junos OS Release 15.1X49-D30, support for the 8-Port Gigabit Ethernet SFP XPIM is restored for SRX550 Service Gateway systems.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1X44-D10 or later for SRX Series Services Gateways.
- Two SRX650 devices connected back-to-back.
- Two 8-port Gigabit Ethernet SFP XPIMs.
- Eight pairs of SFP transceivers as mentioned in 8-Port Gigabit Ethernet SFP XPIM Supported Modules and eight cables to connect them.

Before you begin:

- Establish basic connectivity. See the Getting Started Guide for your device.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 166.

Overview and Topology

In this example, you configure two SRX650 devices. On each device you configure eight interfaces (ge-6/0/0 through ge-6/0/7), set the maximum transmission unit (MTU) value to 9192, and define a logical interface that you can connect to the 8-port SFP XPIM.

Figure 16 on page 232 shows the topology used in this example.

Figure 16: Basic Back-to-Back Device Configuration

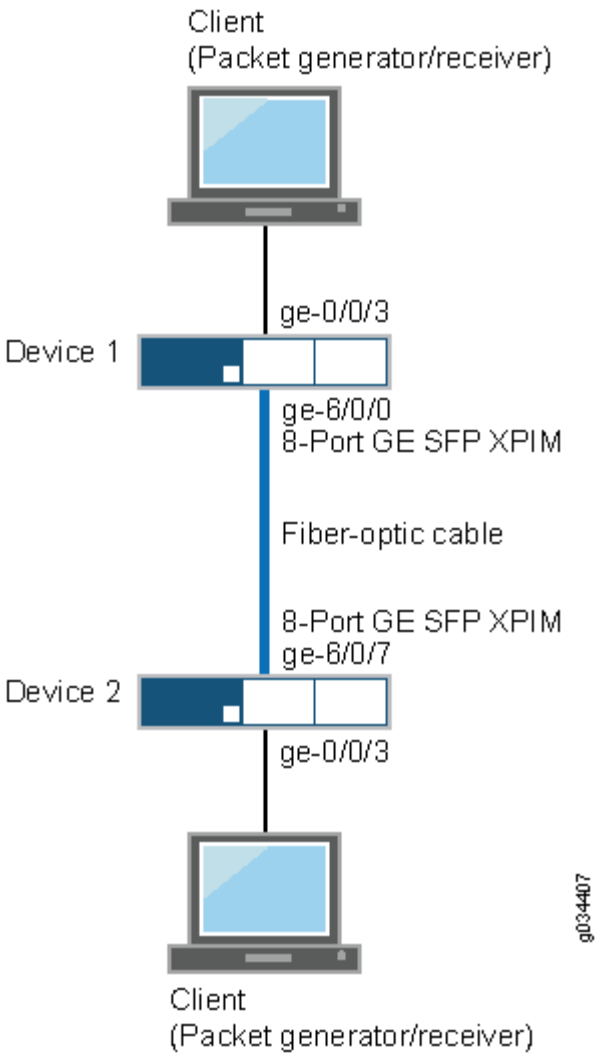

Configuration

IN THIS SECTION

- Procedure | 233

Procedure

CLI Quick Configuration

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

Device 1

```
set interfaces ge-6/0/0 mtu 9192
set interfaces ge-6/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-6/0/1 mtu 9192
set interfaces ge-6/0/1 unit 0 family inet address 11.1.1.1/24
set interfaces ge-6/0/2 mtu 9192
set interfaces ge-6/0/2 unit 0 family inet address 12.1.1.1/24
set interfaces ge-6/0/3 mtu 9192
set interfaces ge-6/0/3 unit 0 family inet address 13.1.1.1/24
set interfaces ge-6/0/4 mtu 9192
set interfaces ge-6/0/4 unit 0 family inet address 14.1.1.1/24
set interfaces ge-6/0/5 mtu 9192
set interfaces ge-6/0/5 unit 0 family inet address 15.1.1.1/24
set interfaces ge-6/0/6 mtu 9192
set interfaces ge-6/0/6 unit 0 family inet address 16.1.1.1/24
set interfaces ge-6/0/7 mtu 9192
set interfaces ge-6/0/7 unit 0 family inet address 17.1.1.1/24
```

Device 2

```
set interfaces ge-6/0/0 mtu 9192
set interfaces ge-6/0/0 unit 0 family inet address 10.1.1.2/24
set interfaces ge-6/0/1 mtu 9192
set interfaces ge-6/0/1 unit 0 family inet address 11.1.1.2/24
set interfaces ge-6/0/2 mtu 9192
set interfaces ge-6/0/2 unit 0 family inet address 12.1.1.2/24
set interfaces ge-6/0/3 mtu 9192
set interfaces ge-6/0/3 unit 0 family inet address 13.1.1.2/24
set interfaces ge-6/0/4 mtu 9192
set interfaces ge-6/0/4 unit 0 family inet address 14.1.1.2/24
set interfaces ge-6/0/5 mtu 9192
set interfaces ge-6/0/5 unit 0 family inet address 15.1.1.2/24
set interfaces ge-6/0/6 mtu 9192
```

```
set interfaces ge-6/0/6 unit 0 family inet address 16.1.1.2/24
set interfaces ge-6/0/7 mtu 9192
set interfaces ge-6/0/7 unit 0 family inet address 17.1.1.2/24
```

Step-by-Step Procedure

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

To configure the interfaces on Device 1:

1. Configure the interface.

```
[edit]
user@host# set interfaces ge-6/0/0
```

2. Assign the maximum transmission unit value for the interface.

```
[edit interfaces ge-6/0/0]
user@host# set mtu 9192
```

3. Add the logical interface.

```
[edit interfaces ge-6/0/0]
user@host# set unit 0 family inet address 10.1.1.1/24
```

NOTE: Repeat these steps for the remaining seven ports on Device 1.

Step-by-Step Procedure

To configure the interfaces on Device 2:

1. Configure the interface.

```
[edit]
user@host# edit interfaces ge-6/0/0
```

2. Assign the maximum transmission unit value for the interface.

```
[edit interfaces ge-6/0/0]
user@host# set mtu 9192
```

3. Add the logical interface.

```
[edit interfaces ge-6/0/0]
user@host# set unit 0 family inet address 10.1.1.2/24
```

NOTE: Repeat these steps for the remaining seven ports on Device 2.

Results

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

Device 1

```
[edit]
user@host# show interfaces
ge-6/0/0 {
    mtu 9192;
    unit 0 {
        family inet {
            address 10.1.1.1/24;
        }
    }
}
ge-6/0/1 {
    mtu 9192;
    unit 0 {
        family inet {
            address 11.1.1.1/24;
        }
    }
}
```

```
ge-6/0/2 {
    mtu 9192;
    unit 0 {
        family inet {
            address 12.1.1.1/24;
        }
    }
}
ge-6/0/3 {
    mtu 9192;
    unit 0 {
        family inet {
            address 13.1.1.1/24;
        }
    }
}
ge-6/0/4 {
    mtu 9192;
    unit 0 {
        family inet {
            address 14.1.1.1/24;
        }
    }
}
ge-6/0/5 {
    mtu 9192;
    unit 0 {
        family inet {
            address 15.1.1.1/24;
        }
    }
}
ge-6/0/6 {
    mtu 9192;
    unit 0 {
        family inet {
            address 16.1.1.1/24;
        }
    }
}
ge-6/0/7 {
    mtu 9192;
    unit 0 {
```

```
        family inet {
            address 17.1.1.1/24;
        }
    }
}
```

Device 2

```
[edit]
user@host# show interfaces
ge-6/0/0 {
    mtu 9192;
    unit 0 {
        family inet {
            address 10.1.1.2/24;
        }
    }
}
ge-6/0/1 {
    mtu 9192;
    unit 0 {
        family inet {
            address 11.1.1.2/24;
        }
    }
}
ge-6/0/2 {
    mtu 9192;
    unit 0 {
        family inet {
            address 12.1.1.2/24;
        }
    }
}
ge-6/0/3 {
    mtu 9192;
    unit 0 {
        family inet {
            address 13.1.1.2/24;
        }
    }
}
```

```
ge-6/0/4 {
    mtu 9192;
    unit 0 {
        family inet {
            address 14.1.1.2/24;
        }
    }
}
ge-6/0/5 {
    mtu 9192;
    unit 0 {
        family inet {
            address 15.1.1.2/24;
        }
    }
}
ge-6/0/6 {
    mtu 9192;
    unit 0 {
        family inet {
            address 16.1.1.2/24;
        }
    }
}
ge-6/0/7 {
    mtu 9192;
    unit 0 {
        family inet {
            address 17.1.1.2/24;
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the Hardware was Properly Installed | **239**
- Verifying the FPC Status | **240**

- Verifying Interface Link Status on Device 1 | **241**
- Verifying the Interface Settings on Device 1 | **242**
- Verifying Interface Link Status on Device 2 | **246**
- Verifying the Interface Settings on Device 2 | **247**

Confirm that the configuration is working properly.

Verifying the Hardware was Properly Installed

Purpose

Verify that the 8-Port Gigabit Ethernet SFP XPIM is installed on the device.

Action

From operational mode, enter the `show chassis hardware` command.

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                AJ3009AA0001      SRX650
Midplane         REV 08   710-023875   AAAK0059
System IO        REV 08   710-023209   AAAJ9290          SRXSME System IO
Routing Engine   REV 13   750-023223   AAAJ1987          RE-SRXSME-SRE6
  ad0    2000 MB  CF 2GB               2009A     0000194075 Compact Flash
  usb0 (addr 1)  DWC OTG root hub 0    vendor 0x0000     uhub0
  usb0 (addr 2)  product 0x005a 90     vendor 0x0409     uhub1
FPC 0                                                    FPC
  PIC 0                                                  4x GE Base PIC
FPC 1            REV 03   750-038290   AADL2016          FPC
FPC 5                                                    FPC
  PIC 0                                                  8x GE SFP gPIM
FPC 6            REV 03   750-037551   AAEC8065          FPC
  PIC 0                                                  8x GE SFP gPIM
    Xcvr 0       REV 01   740-013111   8043353           SFP-T
    Xcvr 1                NON-JNPR     PC602QW           SFP-SX
    Xcvr 2       k        NON-JNPR     BDS3I             SFP-1000BASE-BX10-D
```

```
    Xcvr 3       REV 01   740-011612   9XT702501080      SFP-LH
    Xcvr 4       REV 01   740-011612   9XT702501079      SFP-LH
    Xcvr 5                NON-JNPR     PCH2GTJ           SFP-SX
    Xcvr 6                NON-JNPR     PC604DL           SFP-SX
    Xcvr 7       REV 01   740-011620   5349504           SFP-FX
FPC 8            REV 00   750-038290                     FPC
Power Supply 0
```

Meaning

The output displays the hardware details of the device and a list of all interfaces configured.

Verify that the output contains the following values:

- `FPC 5, PIC 0` —8x SFP gPIM
- `FPC 6, PIC 0` —8x SFP gPIM

NOTE: In the example, the output for 8-Port SFP Gigabit Ethernet XPIM is displayed as 8x GE SFP gPIM.

Verifying the FPC Status

Purpose

Verify that the status of the Flexible PIC Concentrator is online.

Action

From operational mode, enter the `show chassis fpc pic-status` command.

```
user@host>  show chassis fpc pic-status
Slot 0   Online       FPC
  PIC 0  Online       4x GE Base PIC
Slot 1   Present      FPC
Slot 5   Online       FPC
  PIC 0  Online       8x GE SFP gPIM
Slot 6   Online       FPC
```

```
  PIC 0  Online       8x GE SFP gPIM
Slot 8   Present      FPC
```

Meaning

The output shows the FPC status for slot 5 and slot 6 as online. The 8-Port Gigabit Ethernet SFP XPIM is installed in slot 5 and slot 6 of the device.

Verifying Interface Link Status on Device 1

Purpose

Verify that the interface link status is up.

Action

From operational mode, enter the `show interface terse ge-6/0/*` command.

```
user@host> show interface terse ge-6/0/*
```

Output for Device 1

```
Interface               Admin Link Proto    Local                 Remote
ge-6/0/0                up    up
ge-6/0/0.0              up    up   inet     10.1.1.1/24
ge-6/0/1                up    up
ge-6/0/1.0              up    up   inet     11.1.1.1/24
ge-6/0/2                up    up
ge-6/0/2.0              up    up   inet     12.1.1.1/24
ge-6/0/3                up    up
ge-6/0/3.0              up    up   inet     13.1.1.1/24
ge-6/0/4                up    up
ge-6/0/4.0              up    up   inet     14.1.1.1/24
ge-6/0/5                up    up
ge-6/0/5.0              up    up   inet     15.1.1.1/24
ge-6/0/6                up    up
ge-6/0/6.0              up    up   inet     16.1.1.1/24
```

```
ge-6/0/7                up    up
ge-6/0/7.0              up    up   inet     17.1.1.1/24
```

Meaning

The output displays a list of all interfaces configured.

If the link displays `up` for all interfaces, the configuration is working properly. This verifies that the XPIM is up and end-to-end ping is working.

Verifying the Interface Settings on Device 1

Purpose

Verify that the interfaces are configured as expected.

Action

From operational mode, enter the `show interface ge-6/0/0 extensive | no-more` command.

```
user@host>show interface ge-6/0/0 extensive | no-more
```

Output for Device 1

```
Physical interface: ge-6/0/0, Enabled, Physical link is Up
  Interface index: 152, SNMP ifIndex: 544, Generation: 155
  Link-level type: Ethernet, MTU: 9192, 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
  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: 00:26:88:04:0a:a8, Hardware address: 00:26:88:04:0a:a8
  Last flapped   : 2012-07-05 21:58:46 PDT (00:13:29 ago)
  Statistics last cleared: Never
  Traffic statistics:
```

```
 Input  bytes  :                  228                   0 bps
 Output bytes  :                  540                   0 bps
 Input  packets:                    3                   0 pps
 Output packets:                    6                   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
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    3                    3                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                           268              268
  Total packets                            3                3
  Unicast packets                          3                2
  Broadcast packets                        0                1
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
```

```
  Input SA rejects                          0
  Output packet count                                    0
  Output packet pad count                                0
  Output packet error count                              0
  CAM destination filters: 2, CAM source filters: 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: 6
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth              Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95      950000000    95              0      low    none
  3 network-control         5       50000000     5              0      low    none
Interface transmit statistics: Disabled

Logical interface ge-6/0/0.0 (Index 81) (SNMP ifIndex 509) (Generation 146)
  Flags: 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
  Security: Zone: HOST
  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 :       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: 9178, Generation: 162, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
        Generation: 176
```

Meaning

The output displays a list of all interface verification parameters.

Verify the following information in the command output:

- Physical Interface—ge-6/0/0, enabled, physical link is `Up`
- MTU—9192
- Speed—1000 Mbps

If the verification parameters are as expected, the configuration is working properly.

Verifying Interface Link Status on Device 2

Purpose

Verify that the interface link status is up.

Action

From operational mode, enter the `show interface terse ge-6/0/*` command.

```
user@host> show interface terse ge-6/0/*
```

Output for Device 2

```
Interface               Admin Link Proto    Local                 Remote
ge-6/0/0                up    up
ge-6/0/0.0              up    up   inet     10.1.1.2/24
ge-6/0/1                up    up
ge-6/0/1.0              up    up   inet     11.1.1.2/24
ge-6/0/2                up    up
ge-6/0/2.0              up    up   inet     12.1.1.2/24
ge-6/0/3                up    up
ge-6/0/3.0              up    up   inet     13.1.1.2/24
ge-6/0/4                up    up
ge-6/0/4.0              up    up   inet     14.1.1.2/24
ge-6/0/5                up    up
ge-6/0/5.0              up    up   inet     15.1.1.2/24
ge-6/0/6                up    up
ge-6/0/6.0              up    up   inet     16.1.1.2/24
ge-6/0/7                up    up
ge-6/0/7.0              up    up   inet     17.1.1.2/24
```

Meaning

The output displays a list of all interfaces configured.

If the link displays `up` for all interfaces, the configuration is working properly. This verifies that the XPIM is up and end-to-end ping is working.

Verifying the Interface Settings on Device 2

Purpose

Verify that the interfaces are configured as expected.

Action

From operational mode, enter the `show interface ge-6/0/0 extensive | no-more` command.

```
user@host>show interface ge-6/0/0 extensive | no-more
```

Output for Device 2

```
Physical interface: ge-6/0/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 520, Generation: 147
  Link-level type: Ethernet, MTU: 9192, 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
  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: 00:24:dc:17:2f:a8, Hardware address: 00:24:dc:17:2f:a8
  Last flapped   : 2012-07-05 21:59:42 PDT (00:15:32 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 228                 0 bps
   Output bytes  :                 294                 0 bps
   Input  packets:                   3                 0 pps
   Output packets:                   5                 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: 13, 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                    3                    3                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                           268              268
  Total packets                            3                3
  Unicast packets                          2                3
  Broadcast packets                        1                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: 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: 6
CoS information:
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority   Limit
                            %            bps     %           usec
  0 best-effort            95      950000000    95              0      low    none
  3 network-control         5       50000000     5              0      low    none
Interface transmit statistics: Disabled

Logical interface ge-6/0/0.0 (Index 73) (SNMP ifIndex 509) (Generation 146)
  Flags: 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
  Security: Zone: HOST
  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 :       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: 9178, Generation: 162, Route table: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.1.1/24, Local: 10.1.1.2, Broadcast: 10.1.1.255,
      Generation: 176
```

Meaning

The output displays a list of all interface verification parameters.

Verify the following information in the command output:

- Physical Interface—ge-6/0/0, enabled, physical link is Up

- MTU—9192
- Speed—1000 Mbps

If the verification parameters are as expected, the configuration is working properly.

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10, the 8-Port Gigabit Ethernet SFP XPIM is not supported on legacy SRX Series systems.
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10, the 8-Port Gigabit Ethernet SFP XPIM is not supported on legacy SRX Series systems.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

Port Speed on SRX Series Devices

SUMMARY

Learn about port speeds, support for multiple port speeds, and how to configure port speed on SRX Series devices.

IN THIS SECTION

- SRX4600 Port Speed Overview | 251
- Port Speed on SRX5K-IOC4-MRATE | 256

SRX4600 Port Speed Overview

IN THIS SECTION

- Interface Naming Conventions | 253
- Supported Active Physical Ports on SRX4600 to Prevent Oversubscription | 253

Table 27 on page 252 presents the details of SRX4600 port speeds.

Table 27: Port Speed Details and Description

Port Location	Number and Type of Ports	Supported Speeds
FPC0, PIC0 (ports 0–3)	4 chassis cluster ports: • 2 fabric (FAB) • 2 control (CTL)	• 10 Gbps (default) • 1 Gbps (only on CTL ports)
FPC1, PIC0 (ports 0–3)	4 100GbE QSFP28 ports or 40GbE QSFP+ ports	At port or PIC level: • 40 Gbps (default), with QSFP+ optics • 100 Gbps, with QSFP28 optics
FPC1, PIC1 (ports 0–7)	8 10GbE SFP+ ports	• 10 Gbps (default) • 1 Gbps

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

- You need to reboot the chassis cluster for configuration changes (from 10 Gbps to 1 Gbps) to take effect. For more details, see "speed (Chassis Cluster)" on page 1112.
- To configure all 40GbE ports, use the `set chassis fpc 1 pic 0 pic-mode 40G` command.
- To set only the first two 40GbE ports, use the `set chassis fpc 1 pic 0 pic-mode 40G number-of-ports 2` command. This configuration sets only the first two 40GbE ports and disables the last two ports. You need to reboot the device for the configuration to take effect.
- You can channelize each 40GbE port into four 10GbE interfaces by using QSFPP-4X10-GE optics, suitable breakout cables, and the *speed* configuration statement.
- Use the "speed (Gigabit Ethernet interface)" on page 740 configuration to set 1-Gbps speed on PIC 1 ports. 1-Gbps speed is supported only in non-autonegotiation mode. If autonegotiation mode is enabled by default at the remote end, then you must disable it.
- You can configure the interface that is already operating in 10GbE mode to operate in 1GbE mode.
- To prevent oversubscription, configure the number of active ports operating at the configured speed by using the *number-of-ports* statement. The SRX4600 supports a maximum speed of 400 Gbps; the speed cannot be oversubscribed.

- If you try to commit an invalid configuration, the configuration gets committed, but the port is not activated. This is because Junos OS allows you to configure a port before a line card is inserted. You will get an error message in the output of the *show chassis alarms* command and also in the log messages. For example, if you configure four 100GbE interfaces with eight 10GbE interface, then the configuration is invalid.
- The SRX4600 does not support copper SFP transceivers.

For information about interface-naming formats for channelized and nonchannelized interfaces and how to configure SRX Series devices at port level and PIC level, see *Port Speed*.

For information on how to configure the speed at the PIC level, see Table 2 of port speed. For information on how to configure the speed at the port level, see Table 3 of port speed.

For more information about SRX4600 devices, see SRX4600 Services Gateway Hardware Guide.

For information about platforms support, see hardware compatibility tool (HCT).

To view the port speeds on each PIC, execute the *show chassis pic* command.

Interface Naming Conventions

Table 28 on page 253 describes the interface naming convention for a 40GbE interface channelized as four 10GbE interfaces:

Table 28: SRX4600 Interface Naming Convention

Interface Type	Example
4x10GbE	When the 40GbE port et-1/0/0 is channelized into four 10GbE interfaces, the channelized interfaces are named as follows:: xe-1/0/0:0 xe-1/0/0:1 xe-1/0/0:2 xe-1/0/0:3

Supported Active Physical Ports on SRX4600 to Prevent Oversubscription

You can use the *number-of-ports* statement to configure a port as an active port.

Table 29 on page 254 summarizes the SRX4600 active ports with *number-of-ports* and port speed configured at PIC level.

Table 29: SRX4600 Port Speed at PIC level

PIC	Number of Active Ports	Active Port Number at PIC Level		
		10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
PIC 0	0	-	-	-
	1	0	0	0
	2	0, 1	0, 1	0, 1
	3	0, 1, 2	0, 1, 2	0, 1, 2
	4	0, 1, 2, 3	0, 1, 2, 3	0, 1, 2, 3
PIC 1	0	-	-	-
	1	0	-	-
	2	0, 1	-	-
	3	0, 1, 2	-	-
	4	0, 1, 2, 3	-	-
	5	0, 1, 2, 3, 4	-	-
	6	0, 1, 2, 3, 4, 5	-	-
	7	0, 1, 2, 3, 4, 5 ,6	-	-

Table 29: SRX4600 Port Speed at PIC level *(Continued)*

PIC	Number of Active Ports	Active Port Number at PIC Level		
		10-Gigabit Ethernet	40-Gigabit Ethernet	100-Gigabit Ethernet
	8	0, 1, 2, 3, 4, 5, 6, 7	-	-

To prevent oversubscription, you can configure the maximum number of active ports that can operate at the configured speed. Table 30 on page 255 summarizes the maximum number of Gigabit Ethernet ports at PIC and port levels:

Table 30: Maximum Number of Gigabit Ethernet Ports at PIC and Port Level

Port Type	Maximum Number of Ports at PIC Mode (on PIC0 and PIC1)	Maximum Ports Configurable at Port Mode (on PIC0 and PIC1)
10GbE	24 16 ports from PIC 0 and 8 ports from PIC 1.	20 Refers to 12 ports from PIC 0 and 8 ports from PIC 1.
40GbE	4 Only 4 ports from PIC 0. PIC 1 supports only 10-Gbps speed.	4
100GbE	4 Only 4 ports from PIC 0. PIC 1 supports only 10-Gbps speed. **NOTE**: If you configure all four PIC 0 ports as 100GbE interfaces then, the PIC 1 ports are disabled. If you then try to configure any PIC 1 port and commit your configuration, the configuration will be invalid.	4

For information about oversubscription, see Port Speed.

SEE ALSO

Port Speed

speed

show chassis pic

number-of-ports

pic-mode

SRX4600 Services Gateway Hardware Guide

Port Speed on SRX5K-IOC4-MRATE

IN THIS SECTION

- Configuring Port Speed at PIC Level | **256**
- Configuring Port Speed at Port Level | **258**

Each of the twelve ports of PIC 0 and PIC 1 of an SRX5K-IOC4-MRATE supports port speeds of 10 Gbps and 40 Gbps. However, only ports 2 and 5 of both the PICs support port speed of 100 Gbps. You can choose to configure all supported ports of the SRX5K-IOC4-MRATE to operate at the same supported speed or configure all the ports at different supported speeds.

You can configure port speed at the PIC level, to operate all the ports of the SRX5K-IOC4-MRATE at the same speed. That is, you can choose to configure the PIC to operate at a supported speed, and then all the supported ports of the PIC to operate at the configured speed. For example, if you choose to configure PIC 0 at 100-Gbps speed, only ports 2 and 5 of PIC 0 operate at 100-Gbps speed, while the other ports of the PIC are disabled. Similarly, if you choose to configure PIC 0 at 10-Gbps or 40-Gbps speed, all the ports of the PIC are enabled to operate at those speeds. Additionally, you can prevent oversubscription by specifying the number of active physical ports that operate at 10-Gbps, 40-Gbps, and 100-Gbps speeds.

You can configure port speed at the port level, to operate different ports of the SRX5K-IOC4-MRATE at different supported speeds. That is, you configure each port to operate at a supported speed.

The SRX5K-IOC4-MRATE supports an aggregate bandwidth of 480 Gbps, and each of the two PICs supports a bandwidth limit of 240 Gbps. If the aggregate port capacity configured exceeds 240 Gbps per PIC, the configuration is not supported.

Configuring Port Speed at PIC Level

To configure port speed at the PIC level:

1. In configuration mode, navigate to the [edit chassis fpc *fpc-slot* pic *pic-number*] hierarchy level.

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

For example:

```
[edit ]
user@host# edit chassis fpc 4 pic 0
```

2. Configure the pic-mode statement to set the operating speed for the PIC's ports. According to your requirements, you can choose from the options 10G, 40G , or 100G.

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

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set pic-mode 10G
```

3. (Optional) To prevent oversubscription, you can choose to configure the number of ports that operate at the mode configured in Step 2.

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

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set number-of-ports 6
```

4. Verify the configuration.

```
[edit chassis fpc 4 pic 0]
user@host# show
pic-mode 10G;
number-of-ports 6;
```

5. Commit your configuration changes.

If the number-of-ports statement is *not* configured, all the ports that support the speed configured in Step 2 are enabled. That is, depending on that selection, ports 0 through 5 are enabled for speeds of 10-gigabit or 40-gigabit, while ports 2 and 5 are enabled for 100-gigabit. You can also use the number-of-ports statement to disable certain ports. Table 1 below, lists the physical ports that are enabled when the number-of-ports statement is configured.

Table 31: Active Physical Ports on SRX5K-IOC4-MRATE Based on the number-of-ports Configuration

Ports Configured (number-of-ports Statement)	Active Physical Ports for Different Configured Speeds		
	10-Gigabit	40-Gigabit	100-Gigabit
1	0	0	2
2	0, 1	0, 1	2, 5
3	0, 1, 2	0, 1, 2	2, 5
4	0, 1, 2, 3	0, 1, 2, 3	2, 5
5	0, 1, 2, 3, 4	0, 1, 2, 3, 4	2, 5
6	0, 1, 2, 3, 4, 5	0, 1, 2, 3, 4, 5	2, 5

Configuring Port Speed at Port Level

To configure port speed at the port level:

1. In configuration mode, navigate to the [edit chassis fpc *fpc-slot* pic *pic-number*] hierarchy level.

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

For example:

```
[edit]
user@host# edit chassis fpc 4 pic 0
```

2. To indicate the speed at which the ports operate, configure the `speed` statement for the desired ports. According to your requirements, you can choose the `10g`, `40g`, or `100g` speed options.

```
[edit chassis fpc fpc-slot pic pic-number]
user@host# set port port-number speed (10g | 40g | 100g)
```

For example:

```
[edit chassis fpc 4 pic 0]
user@host# set port 0 speed 10g
user@host# set port 1 speed 10g
user@host# set port 2 speed 100g
user@host# set port 3 speed 40g
```

All the twelve ports of PIC 0 and PIC 1 of an SRX5K-IOC4-MRATE support 10-Gbps and 40-Gbps port speeds. However, only ports 2 and 5 of both the PICs support 100-Gbps speed.

3. Verify the configuration.

```
[edit chassis fpc 4 pic 0]
user@host# show
port 0 {
    speed 10g;
}
port 1 {
    speed 10g;
}
port 2 {
    speed 100g;
}
port 3 {
    speed 40g;
}
```

You can verify the 40-Gbps and 100-Gbps ports configured as 10-Gbps by using the `show interfaces terse` command.

```
[edit chassis fpc 10 pic 0 port 0]
+     speed 10g;

user@host# show interfaces terse
..
xe-10/0/0:0             up    down
xe-10/0/0:1             up    down
xe-10/0/0:2             up    down
xe-10/0/0:3             up    down
```

4. Commit your configuration changes.

Note the following when configuring port speed on an SRX5K-IOC4-MRATE:

- If port speed is not configured, all ports of the SRX5K-IOC4-MRATE operate as four 10-Gigabit Ethernet interfaces by default. Therefore, when booting the MPC:
  - If port speed is not configured or if invalid port speeds are configured, each port operates as four 10-Gigabit Ethernet interfaces. An alarm is generated to indicate that the ports of the SRX5K-IOC4-MRATE are operating as four 10-Gigabit Ethernet interfaces.
  - If valid port speeds are configured, the MPC PICs operate at the configured speed.
- • When you change an existing port speed configuration at the port level, you must reset the SRX5K-IOC4-MRATE PIC for the configuration to take effect. An alarm is generated indicating the change in port speed configuration.
- • When you change an existing port speed configuration with an *invalid* port speed configuration, an alarm is generated indicating that the port speed configuration is invalid. The MPC continues to operate using the previously configured valid port speed configuration. However, if the MPC or PIC is restarted with the committed invalid port configuration, all ports of the MPC operate as four 10-Gigabit Ethernet interfaces by default.
- • You cannot configure port speed at the PIC level and the port level simultaneously. Error messages are displayed when you try to commit such configurations.
- • When you configure port speed at the port level, only the configured ports are enabled. Other ports are disabled.
- Logical interfaces can be created only on ports that are enabled.
- You must restart the chassis when you change the port profile configuration.

Power over Ethernet

IN THIS SECTION

- Power over Ethernet Overview | 261
- Example: Configure PoE Interface | 268

Power over Ethernet (PoE) is the implementation of the IEEE 802.3 AF and IEEE 802.3 AT standards that allow both data and electrical power to pass over a copper Ethernet LAN cable. The topics below discuss the overview and configuration details of PoE, and disabling a PoE interface on security devices.

Power over Ethernet Overview

IN THIS SECTION

- SRX Series Services Gateway PoE Specifications | 261
- PoE Classes and Power Ratings | 267
- PoE Options | 267

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

You can configure the SRX Series device to act as power sourcing equipment (PSE), supplying power to powered devices that are connected on designated ports. For more information about PoE, see Power over Ethernet (PoE) User Guide for EX Series Switches.

This topic contains the following sections:

SRX Series Services Gateway PoE Specifications

Table 32 on page 262 lists the PoE specifications for the SRX210, SRX220, SRX240, SRX320, SRX650, and SRX550 M devices. (Platform support depends on the Junos OS release in your installation.)

Table 32: PoE Specifications for the SRX210, SRX220, SRX240, SRX320, and SRX650 Devices

Specifications	For SRX210 Device	For SRX220 Device	For SRX240 Device	For SRX320 PoE Device		For SRX 550 M device	For SRX650 Device
Supported standards	• IEEE 802.3 AF • Legacy (pre-standards)	• IEEE 802.3 AF • IEEE 802.3 AT (PoE +) • Legacy (pre-standards)	• IEEE 802.3 AF • IEEE 802.3 AT (PoE +) • Legacy (pre-standards)	• IEEE 802.3 AF • IEEE 802.3 AT (PoE) • Legacy (pre-standards)		• IEEE 802.3 AF • IEEE 802.3 AT (PoE +) • Legacy (pre-standards)	• IEEE 802.3 AF • IEEE 802.3 AT (PoE +) • Legacy (pre-standards)

Table 32: PoE Specifications for the SRX210, SRX220, SRX240, SRX320, and SRX650 Devices *(Continued)*

Specifications	For SRX210 Device	For SRX220 Device	For SRX240 Device	For SRX320 PoE Device		For SRX 550 M device	For SRX650 Device

Supported ports	Supported on two Gigabit Ethernet ports and two Fast Ethernet ports (ge-0/0/0, ge-0/0/1, fe-0/0/2, and fe-0/0/3).	Supported on all 8 Gigabit Ethernet ports (ge-0/0/0 to ge-0/0/7).	Supported on all 16 Gigabit Ethernet ports (ge-0/0/0 to ge-0/0/15).	Supported on all 6 Copper (RJ45) Gigabit Ethernet ports (ge-0/0/0 to ge-0/0/5).		Supported on 16GE-POE xPIM card	Supported on the following ports:
							• Slot 2 or 6 on 16 Gigabit Ethernet ports
							  • ge-2/0/0 to ge-2/0/15
							  • ge-6/0/0 to ge-6/0/15
							• Slot 2 or 6 on 24 Gigabit Ethernet ports
							  • ge-2/0/0 to ge-2/0/23

Table 32: PoE Specifications for the SRX210, SRX220, SRX240, SRX320, and SRX650 Devices *(Continued)*

Specifications	For SRX210 Device	For SRX220 Device	For SRX240 Device	For SRX320 PoE Device		For SRX 550 M device	For SRX650 Device
							• ge-6/0/0 to ge-6/0/23

Table 32: PoE Specifications for the SRX210, SRX220, SRX240, SRX320, and SRX650 Devices *(Continued)*

Specifications	For SRX210 Device	For SRX220 Device	For SRX240 Device	For SRX320 PoE Device		For SRX 550 M device	For SRX650 Device
Total PoE power sourcing capacity	50 W	120 W	150 W	180 W		The 645 watts AC and 645 watts DC power supplies support the following capacities: • 250 watts on a single power supply, or with redundancy using the two-power-supply option. • 500 watts with the two-power-supply option operating as nonredundant.	The 645 watts AC and 645 watts DC power supplies support the following capacities: • 250 watts on a single power supply, or with redundancy using the two-power-supply option. • 500 watts with the two-power-supply option operating as nonredundant.

Table 32: PoE Specifications for the SRX210, SRX220, SRX240, SRX320, and SRX650 Devices *(Continued)*

Specifications	For SRX210 Device	For SRX220 Device	For SRX240 Device	For SRX320 PoE Device		For SRX 550 M device	For SRX650 Device
Default per port power limit	15.4 W	15.4 W	15.4 W	30 W		15.4 W	15.4 W
Maximum per port power limit	30 W	30W	30 W	30 W		30 W	30 W
Power management modes	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.		• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.	• Static: Power allocated for each interface can be configured. • Class: Power allocated for interfaces is based on the class of powered device connected.

PoE Classes and Power Ratings

Table 33 on page 267 lists the classes and their power ratings as specified by the IEEE standards.

Table 33: SRX Series Devices PoE Specifications

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

PoE Options

When you configure PoE, you must enable the PoE interface for the port to provide power to a connected, powered device. In addition, you can configure the following PoE features:

- Port priority—Sets port priority. When it is not possible to maintain power to all connected ports, lower priority ports are powered off before higher priority ports. When you connect a device on a higher-priority port, a lower priority port will be powered off automatically if available power is insufficient to power on the higher priority port. (For the ports with the same priority configuration, ports on the left are given higher priority than the ports on the right.)
- Maximum available wattage power available to a port—Sets the maximum amount of power that can be supplied to the port. Default wattage per port is 15.4 watts.
- PoE power consumption logging—Allows logging of per-port PoE power consumption. The telemetries section is disabled by default and must be explicitly specified to enable logging. Default telemetry duration is 1 hour, and the interval is 5 minutes.
- PoE power management mode—Has two modes:

- Class—Power is allocated dynamically using the classification process.
- Static—Power is allocated based on the maximum power configuration.

- Reserve power—Specified amount of power is reserved for the gateway in case of a spike in PoE consumption. The default is 0.

Example: Configure PoE Interface

IN THIS SECTION

- Verification | 271

In this topic you can learn to configure the PoE interface on all interfaces, an individual interface, and how to disable a PoE interface. Below table specifies the CLI quick configuration commands used for configuring PoE interfaces.

CLI Quick Configuration

Use the below table to view the CLI quick configuration commands to configure PoE on individual and all interfaces, and also to disabe the interface.

Table 34: CLI Quick Configuration

Configuration Step	CLI Quick Configuration Commands
Configure PoE on an individual interface.	`set poe interface ge-0/0/0 priority high maximum-power 15.4 telemetries` `set poe management static guard-band 15`
Configure PoE on all individual interfaces.	`set poe interface all priority low maximum-power 15.4 telemetries` `set poe management static guard-band 15`

Configure PoE Interfaces

Below table describes the step-by-step to configure PoE interfaces on your security device.

Table 35: PoE Interfaces Configuration

Configuration Step	Command

Step 1: Enable PoE	For an individual interface:
	
	`[edit`
	`user@host# edit poe interface ge-0/0/0`
	
	For all interfaces:
	
	`[edit]`
	`user@host# edit poe interface all`

Step 2: Set the power port priority.	For an individual interface:
	
	`[edit poe interface ge-0/0/0]`
	`user@host# set priority high`
	
	For all interfaces:
	
	`[edit poe interface all]`
	`user@host# set priority low`

Step 3: Set the maximum PoE wattage available for a port.	For an individual interface:
	
	`[edit poe interface ge-0/0/0]`
	`user@host# set maximum power 15.4`
	
	For all interfaces:
	
	`[edit poe interface all]`
	`user@host# set maximum power 15.4`

Table 35: PoE Interfaces Configuration *(Continued)*

Configuration Step	Command

Step 4: Enable logging of PoE power consumption.	For an individual interface:
	
	`[edit poe interface ge-0/0/0]`
	`user@host# set telemetries`
	
	For all interfaces:
	
	`[edit poe interface all]`
	`user@host# set telemetries`

Set the PoE management mode.	`[edit]` `user@host# set poe management static`
Step 6: Reserve power wattage in case of a spike in PoE consumption.	`[edit]` `user@host# set poe guard-band 15`
Step 7: (Optional) Disable PoE on all interfaces.	`[edit]` `user@host# set poe interface all disable`
Step 8: (Optional) Disable PoE on a specific interface.	`[edit]` `user@host# set poe interface ge-0/0/0 disable`
Step 9: If you are done configuring the device, commit the configuration.	`[edit]` `user@host# commit`

Use the `show poe interface ge-0/0/0` and `show poe interface all` command to see the output of the configuration. To verify the configuration is working properly, enter the `show poe interface`command.

Verification

Purpose

Verify that PoE interface is enabled on individual and all interfaces, also check how to disable PoE interface. (The device used in this example is the SRX240 or SRX340 Services Gateway, depending on the Junos OS release in the installation.)

Action

- To display information about the parameters configured on PoE interface.

```
user@host> show poe interface ge-0/0/1
PoE interface status:
PoE interface                : ge-0/0/1
Administrative status        : Enabled
Operational status           : Powered-up
Power limit on the interface : 15.4 W
Priority                     : High
Power consumed               : 6.6 W
Class of power device        : 0
```

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

  For all records:

```
user@host> show poe telemetries interface ge-0/0/1 all
Sl No Timestamp Power Voltage
1 Fri Jan 04 11:41:15 2009 5.1 W 47.3 V
2 Fri Jan 04 11:40:15 2009 5.1 W 47.3 V
3 Fri Jan 04 11:39:15 2009 5.1 W 47.3 V
4 Fri Jan 04 11:38:15 2009 0.0 W 0.0 V
5 Fri Jan 04 11:37:15 2009 0.0 W 0.0 V
6 Fri Jan 04 11:36:15 2009 6.6 W 47.2 V
7 Fri Jan 04 11:35:15 2009 6.6 W 47.2 V
```

  For a specific number of records:

```
user@host> show poe telemetries interface ge-0/0/1 5
Sl No Timestamp Power Voltage
1 Fri Jan 04 11:31:15 2009 6.6 W 47.2 V
```

```
2 Fri Jan 04 11:30:15 2009 6.6 W 47.2 V
3 Fri Jan 04 11:29:15 2009 6.6 W 47.2 V
4 Fri Jan 04 11:28:15 2009 6.6 W 47.2 V
5 Fri Jan 04 11:27:15 2009 6.6 W 47.2 V
```

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

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

```
user@host> show poe controller
Controller  Maximum   Power          Guard band  Management
index       power     consumption
   0        150.0 W    0.0 W         0 W         Static
```

- Verify that the PoE interfaces on the device are enabled and set to the desired priority settings. (The device used here is the SRX340 Services Gateway.)

```
user@host> show poe interface all
```

```
Interface Admin status Oper status Max power Priority Power consumption Class
 ge-0/0/0 Enabled      Searching   15.4W     Low      0.0W                0
 ge-0/0/1 Enabled      Powered-up  15.4W     High     6.6W                0
 ge-0/0/2 Disabled     Disabled    15.4W     Low      0.0W                0
 ge-0/0/3 Disabled     Disabled    15.4W     Low      0.0W                0
```

This output shows that the device has four PoE interfaces of which two are enabled with default values. One port has a device connected that is drawing power within expected limits.

5

CHAPTER

Configuring Interface Encapsulation

Interface Encapsulation Overview | 274

Configuring GRE Keepalive Time | 281

Configuring Point-to-Point Protocol over Ethernet | 307

Interface Encapsulation Overview

IN THIS SECTION

- Understanding Physical Encapsulation on an Interface | 274
- Understanding Frame Relay Encapsulation on an Interface | 275
- Understanding Point-to-Point Protocol | 277
- Understanding High-Level Data Link Control | 280

The below topics discuss the overview of overview of physical encapsulation, frame relay encapsulation, point-to-point protocol and high-level data link control.

Understanding Physical Encapsulation on an Interface

Encapsulation is the process by which a lower level protocol accepts a message from a higher level protocol and places it in the data portion of the lower level frame. As a result, datagrams transmitted through a physical network have a sequence of headers: the first header for the physical network (or Data Link Layer) protocol, the second header for the Network Layer protocol (IP, for example), the third header for the Transport Layer protocol, and so on.

The following encapsulation protocols are supported on physical interfaces:

- Frame Relay Encapsulation. See "Understanding Frame Relay Encapsulation on an Interface" on page 275.
- Point-to-Point Protocol. See "Understanding Point-to-Point Protocol" on page 277.
- Point-to-Point Protocol over Ethernet. See "Understanding Point-to-Point Protocol over Ethernet" on page 308.
- High-Level Data Link Control. See "Understanding High-Level Data Link Control" on page 280.

SEE ALSO

Understanding Interfaces | 2

Understanding Frame Relay Encapsulation on an Interface

IN THIS SECTION

- Virtual Circuits | 276
- Switched and Permanent Virtual Circuits | 276
- Data-Link Connection Identifiers | 276
- Congestion Control and Discard Eligibility | 276

The Frame Relay packet-switching protocol operates at the Physical Layer and Data Link Layer in a network to optimize packet transmissions by creating virtual circuits between hosts. Figure 17 on page 275 shows a typical Frame Relay network.

Figure 17: Frame Relay Network

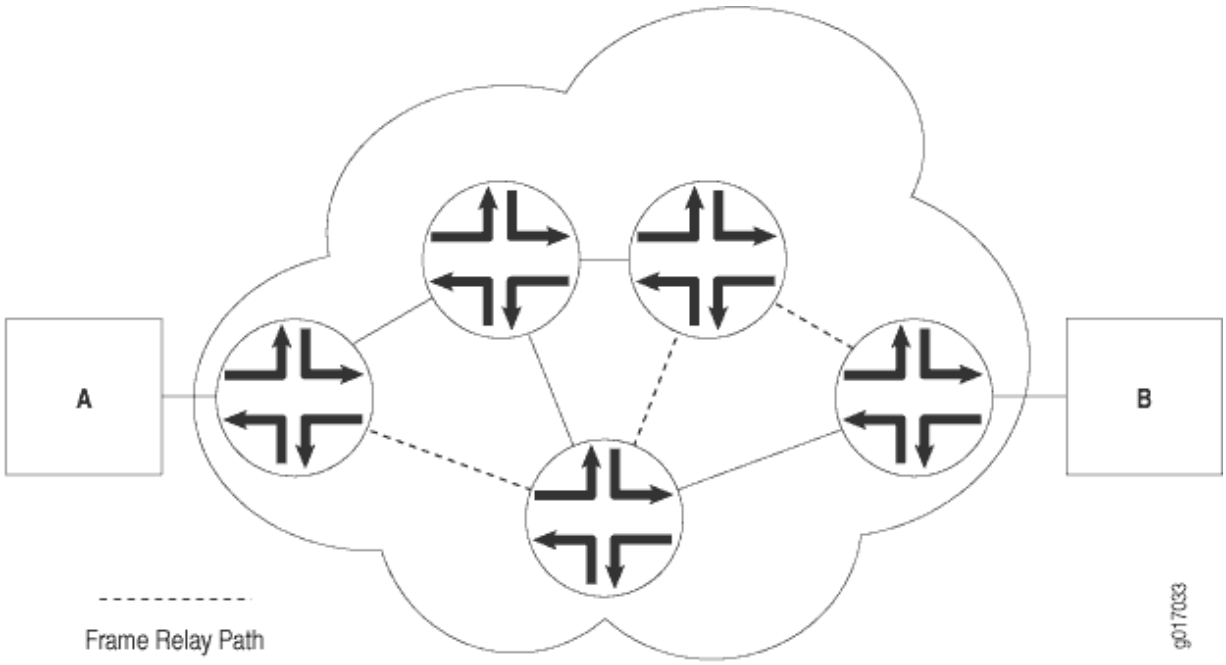

Figure 17 on page 275 shows multiple paths from Host A to Host B. In a typical routed network, traffic is sent from device to device with each device making routing decisions based on its own routing table. In a packet-switched network, the paths are predefined. Devices switch a packet through the network according to predetermined next-hops established when the virtual circuit is set up.

This topic contains the following sections:

Virtual Circuits

A virtual circuit is a bidirectional path between two hosts in a network. Frame Relay virtual circuits are logical connections between two hosts that are established either by a call setup mechanism or by an explicit configuration.

A virtual circuit created through a call setup mechanism is known as a switched virtual circuit (SVC). A virtual circuit created through an explicit configuration is called a permanent virtual circuit (PVC).

Switched and Permanent Virtual Circuits

Before data can be transmitted across an SVC, a signaling protocol like ISDN must set up a call by the exchange of setup messages across the network. When a connection is established, data is transmitted across the SVC. After data transmission, the circuit is torn down and the connection is lost. For additional traffic to pass between the same two hosts, a subsequent SVC must be established, maintained, and terminated.

Because PVCs are explicitly configured, they do not require the setup and teardown of SVCs. Data can be switched across the PVC whenever a host is ready to transmit. SVCs are useful in networks where data transmission is sporadic and a permanent circuit is not needed.

Data-Link Connection Identifiers

An established virtual circuit is identified by a data-link connection identifier (DLCI). The DLCI is a value from 16 through 1022. (Values 1 through 15 are reserved.) The DLCI uniquely identifies a virtual circuit locally so that devices can switch packets to the appropriate next-hop address in the circuit. Multiple paths that pass through the same transit devices have different DLCIs and associated next-hop addresses.

Congestion Control and Discard Eligibility

Frame Relay uses the following types of congestion notification to control traffic within a Frame Relay network. Both are controlled by a single bit in the Frame Relay header.

- Forward explicit congestion notification (FECN)
- Backward explicit congestion notification (BECN)

Traffic congestion is typically defined in the buffer queues on a device. When the queues reach a predefined level of saturation, traffic is determined to be congested. When traffic congestion occurs in a virtual circuit, the device experiencing congestion sets the congestion bits in the Frame Relay header to

1. As a result, transmitted traffic has the FECN bit set to 1, and return traffic on the same virtual circuit has the BECN bit set to 1.

When the FECN and BECN bits are set to 1, they provide a congestion notification to the source and destination devices. The devices can respond in either of two ways: to control traffic on the circuit by sending it through other routes, or to reduce the load on the circuit by discarding packets.

If devices discard packets as a means of congestion (flow) control, Frame Relay uses the discard eligibility (DE) bit to give preference to some packets in discard decisions. A DE value of 1 indicates that the frame is of lower importance than other frames and more likely to be dropped during congestion. Critical data (such as signaling protocol messages) without the DE bit set is less likely to be dropped.

Understanding Point-to-Point Protocol

IN THIS SECTION

- Link Control Protocol | **278**
- PPP Authentication | **278**
- Network Control Protocols | **279**
- Magic Numbers | **279**
- CSU/DSU Devices | **280**

The Point-to-Point Protocol (PPP) is an encapsulation protocol for transporting IP traffic across point-to-point links. PPP is made up of three primary components:

- Link Control Protocol (LCP)—Establishes working connections between two points.
- Authentication protocol—Enables secure connections between two points.
- Network control protocol (NCP)—Initializes the PPP protocol stack to handle multiple Network Layer protocols, such as IPv4, IPv6, and Connectionless Network Protocol (CLNP).

This topic contains the following sections:

Link Control Protocol

LCP is responsible for establishing, maintaining, and tearing down a connection between two endpoints. LCP also tests the link and determines whether it is active. LCP establishes a point-to-point connection as follows:

1. LCP must first detect a clocking signal on each endpoint. However, because the clocking signal can be generated by a network clock and shared with devices on the network, the presence of a clocking signal is only a preliminary indication that the link might be functioning.
2. When a clocking signal is detected, a PPP host begins transmitting PPP Configure-Request packets.
3. If the remote endpoint on the point-to-point link receives the Configure-Request packet, it transmits a Configure-Acknowledgement packet to the source of the request.
4. After receiving the acknowledgement, the initiating endpoint identifies the link as established. At the same time, the remote endpoint sends its own request packets and processes the acknowledgement packets. In a functioning network, both endpoints treat the connection as established.

During connection establishment, LCP also negotiates connection parameters such as FCS and HDLC framing. By default, PPP uses a 16-bit FCS, but you can configure PPP to use either a 32-bit FCS or a 0-bit FCS (no FCS). Alternatively, you can enable HDLC encapsulation across the PPP connection.

After a connection is established, PPP hosts generate Echo-Request and Echo-Response packets to maintain a PPP link.

PPP Authentication

PPP's authentication layer uses a protocol to help ensure that the endpoint of a PPP link is a valid device. Authentication protocols include the Password Authentication Protocol (PAP), the Extensible Authentication Protocol (EAP), and the Challenge Handshake Authentication Protocol (CHAP). CHAP is the most commonly used.

NOTE: Support for user id and the password to comply with full ASCII character set is supported through RFC 2486.

The user can enable or disable the RFC 2486 support under the PPP options. The RFC 2486 is disabled by default, and enable the support globally use the command `set access ppp-options compliance rfc 2486`".

CHAP ensures secure connections across PPP links. After a PPP link is established by LCP, the PPP hosts at either end of the link initiate a three-way CHAP handshake. Two separate CHAP handshakes are required before both sides identify the PPP link as established.

CHAP configuration requires each endpoint on a PPP link to use a shared secret (password) to authenticate challenges. The shared secret is never transmitted over the wire. Instead, the hosts on the PPP connection exchange information that enables both to determine that they share the same secret. Challenges consist of a hash function calculated from the secret, a numeric identifier, and a randomly chosen challenge value that changes with each challenge. If the response value matches the challenge value, authentication is successful. Because the secret is never transmitted and is required to calculate the challenge response, CHAP is considered very secure.

PAP authentication protocol uses a simple two-way handshake to establish identity. PAP is used after the link establishment phase (LCP up), during the authentication phase. Junos OS can support PAP in one direction (egress or ingress), and CHAP in the other.

Network Control Protocols

After authentication is completed, the PPP connection is fully established. At this point, any higher level protocols (for example, IP protocols) can initialize and perform their own negotiations and authentication.

PPP NCPs include support for the following protocols. IPCP and IPv6CP are the most widely used on SRX Series devices.

- IPCP—IP Control Protocol
- IPv6CP—IPv6 Control Protocol
- OSINLCP—OSI Network Layer Control Protocol (includes IS-IS, ES-IS, CLNP, and IDRP)

Magic Numbers

Hosts running PPP can create “magic” numbers for diagnosing the health of a connection. A PPP host generates a random 32-bit number and sends it to the remote endpoint during LCP negotiation and echo exchanges.

In a typical network, each host's magic number is different. A magic number mismatch in an LCP message informs a host that the connection is not in loopback mode and traffic is being exchanged bidirectionally. If the magic number in the LCP message is the same as the configured magic number, the host determines that the connection is in loopback mode, with traffic looped back to the transmitting host.

Looping traffic back to the originating host is a valuable way to diagnose network health between the host and the loopback location. To enable loopback testing, telecommunications equipment typically supports channel service unit/data service unit (CSU/DSU) devices.

CSU/DSU Devices

A channel service unit (CSU) connects a terminal to a digital line. A data service unit (DSU) performs protective and diagnostic functions for a telecommunications line. Typically, the two devices are packaged as a single unit. A CSU/DSU device is required for both ends of a T1 or T3 connection, and the units at both ends must be set to the same communications standard.

A CSU/DSU device enables frames sent along a link to be looped back to the originating host. Receipt of the transmitted frames indicates that the link is functioning correctly up to the point of loopback. By configuring CSU/DSU devices to loop back at different points in a connection, network operators can diagnose and troubleshoot individual segments in a circuit.

Understanding High-Level Data Link Control

IN THIS SECTION

- HDLC Stations | **280**
- HDLC Operational Modes | **281**

High-Level Data Link Control (HDLC) is a bit-oriented, switched and nonswitched link-layer protocol. HDLC is widely used because it supports half-duplex and full-duplex connections, point-to-point and point-to-multipoint networks, and switched and nonswitched channels.

This topic contains the following sections:

HDLC Stations

Nodes within a network running HDLC are called stations. HDLC supports three types of stations for data link control:

- Primary stations—Responsible for controlling the secondary and combined other stations on the link. Depending on the HDLC mode, the primary station is responsible for issuing acknowledgement packets to allow data transmission from secondary stations.
- Secondary stations—Controlled by the primary station. Under normal circumstances, secondary stations cannot control data transmission across the link with the primary station, are active only when requested by the primary station, and can respond to the primary station only (not to other secondary stations). All secondary station frames are response frames.

- Combined stations—A combination of primary and secondary stations. On an HDLC link, all combined stations can send and receive commands and responses without any permission from any other stations on the link and cannot be controlled by any other station.

HDLC Operational Modes

HDLC runs in three separate modes:

- Normal Response Mode (NRM)—The primary station on the HDLC link initiates all information transfers with secondary stations. A secondary station on the link can transmit a response of one or more information frames only when it receives explicit permission from the primary station. When the last frame is transmitted, the secondary station must wait for explicit permission before it can transmit more frames.

  NRM is used most widely for point-to-multipoint links, in which a single primary station controls many secondary stations.

- Asynchronous Response Mode (ARM)—The secondary station can transmit either data or control traffic at any time, without explicit permission from the primary station. The primary station is responsible for error recovery and link setup, but the secondary station can transmit information at any time.

  ARM is used most commonly with point-to-point links, because it reduces the overhead on the link by eliminating the need for control packets.

- Asynchronous Balance Mode (ABM)—All stations are combined stations. Because no other station can control a combined station, all stations can transmit information without explicit permission from any other station. ABM is not a widely used HDLC mode.

Configuring GRE Keepalive Time

IN THIS SECTION

- Understanding GRE Keepalive Time | **282**
- Configuring GRE Keepalive Time | **283**
- Example: GRE Configuration | **288**
- Example: Configuring GRE over IPsec Tunnels | **295**
- Example: Configuring a GRE Tunnel When the Tunnel Destination Is in a Routing Instance | **300**

Generic routing encapsulation (GRE) tunnel interfaces do not have a built-in mechanism for detecting when a tunnel is down. Keepalive messages help the GRE tunnel interfaces to detect when a tunnel is down. The topics below discuss the working and configuration of GRE keepalive time.

Understanding GRE Keepalive Time

Generic routing encapsulation (GRE) tunnel interfaces do not have a built-in mechanism for detecting when a tunnel is down. You can enable keepalive messages to serve as the detection mechanism.

Keepalive times are only configurable for the ATM-over-ADSL interface, which is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM starting in Junos OS Release 15.1X49-D10. Keepalive times are enabled by default for other interfaces.

Keepalives can be configured on the physical or on the *logical interface*. If configured on the physical interface, keepalives are sent on all logical interfaces that are part of the physical interface. If configured on a individual logical interface, keepalives are only sent to that logical interface. In addition to configuring a keepalive, you must configure the hold time.

You can configure the keepalives on a generic routing encapsulation (GRE) tunnel interface by including both the `keepalive-time` statement and the `hold-time` statement at the `[edit protocols oam gre-tunnel interface` *`interface-name`*`]` hierarchy level.

NOTE: For proper operation of keepalives on a GRE interface, you must also include the `family inet` statement at the `[edit interfaces` *`interface-name`* `unit` *`unit`*`]` hierarchy level. If you do not include this statement, the interface is marked as down.

SEE ALSO

keepalive-time

hold-time

Configuring GRE Keepalive Time

IN THIS SECTION

- Configuring Keepalive Time and Hold time for a GRE Tunnel Interface | 283
- Display GRE Keepalive Time Configuration | 284
- Display Keepalive Time Information on a GRE Tunnel Interface | 285

Keepalive times are only configurable for the ATM-over-ADSL interface, which is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM starting in Junos OS Release 15.1X49-D10.

Configuring Keepalive Time and Hold time for a GRE Tunnel Interface

You can configure the keepalives on a generic routing encapsulation (GRE) tunnel interface by including both the `keepalive-time` statement and the `hold-time` statement at the `[edit protocols oam gre-tunnel interface` *`interface-name`*`]` hierarchy level.

NOTE: For proper operation of keepalives on a GRE interface, you must also include the `family inet` statement at the `[edit interfaces` *`interface-name`* `unit` *`unit`*`]` hierarchy level. If you do not include this statement, the interface is marked as down.

To configure a GRE tunnel interface:

1. Configure the GRE tunnel interface at `[edit interfaces` *`interface-name`* `unit` *`unit-number`*`]` hierarchy level, where the interface name is gr-x/y/z, and the family is set as `inet`.

```
user@host# set interfaces interface-name unit unit-number family family-name
```

2. Configure the rest of the GRE tunnel interface options based on requirement.

To configure keepalive time for a GRE tunnel interface:

1. Configure the Operation, Administration, and Maintenance (OAM) protocol at the `[edit protocols]` hierarchy level for the GRE tunnel interface.

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

2. Configure the GRE tunnel interface option for OAM protocol.

```
[edit protocols oam]
user@host# edit gre-tunnel interface interface-name
```

3. Configure the keepalive time from 1 through 50 seconds for the GRE tunnel interface.

```
[edit protocols oam gre-tunnel interface interface-name]
user@host# set keepalive-time seconds
```

4. Configure the hold time from 5 through 250 seconds. Note that the hold time must be at least twice the keepalive time.

```
[edit protocols oam gre-tunnel interface interface-name]
user@host# set hold-time seconds
```

Display GRE Keepalive Time Configuration

IN THIS SECTION

- Purpose | 284
- Action | 285

Purpose

Display the configured keepalive time value as 10 and hold time value as 30 on a GRE tunnel interface (for example, gr-1/1/10.1):

Action

To display the configured values on the GRE tunnel interface, run the `show oam gre-tunnel` command at the `[edit protocols]` hierarchy level:

```
[edit protocols]
user@host# show oam gre-tunnel
    interface gr-1/1/10.1
{

        keepalive-time
10;

        hold-time
30;

    }
```

Display Keepalive Time Information on a GRE Tunnel Interface

IN THIS SECTION

- Purpose | 285
- Action | 285
- Meaning | 287

Purpose

Display the current status information of a GRE tunnel interface when keepalive time and hold time parameters are configured on it and when the hold time expires.

Action

To verify the current status information on a GRE tunnel interface (for example, gr-3/3/0.3), run the `show interfaces gr-3/3/0.3 terse` and `show interfaces gr-3/3/0.3 extensive` operational commands.

show interfaces gr-3/3/0.3 terse

```
user@host> show interfaces gr-3/3/0.3 terse


                    Interface               Admin Link Proto    Local                 Remote

                    gr-3/3/0.3              up    up   inet     200.1.3.1/24

                    mpls
```

show interfaces gr-3/3/0.3 extensive

```
user@host> show interfaces gr-3/3/0.3 extensive
Logical interface gr-3/3/0.3 (Index 73) (SNMP ifIndex 594) (Generation 900)

                        Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header
10.1.19.11:10.1.19.12:47:df:64:0000000000000000 Encapsulation: GRE-NULL

                        Gre keepalives configured: On, Gre keepalives adjacency state: down


^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^

                          Traffic statistics:

     Input  bytes  :             15629992

     Output bytes  :             15912273

     Input  packets:               243813

     Output packets:               179476

     Local statistics:

     Input  bytes  :             15322586

     Output bytes  :             15621359

     Input  packets:               238890
```

```
Output packets:              174767

Transit statistics:

Input  bytes  :              307406                   0 bps

Output bytes  :              290914                   0 bps

Input  packets:                4923                   0 pps

Output packets:                4709                   0 pps

Protocol inet, MTU: 1476, Generation: 1564, Route table: 0

Flags: Sendbcast-pkt-to-re

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

  ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^

  Destination: 200.1.3/24, Local: 200.1.3.1, Broadcast: 200.1.3.255, Generation: 1366

  Protocol mpls, MTU: 1464, Maximum labels: 3, Generation: 1565, Route table: 0
```

NOTE: When the hold time expires:

- The GRE tunnel will stay up even though the interface cannot send or receive traffic.
- The `Link` status will be `Up` and the `Gre keepalives adjacency state` will be `Down`.

Meaning

The current status information of a GRE tunnel interface with keepalive time and hold time parameters is displayed as expected when the hold time expires.

Example: GRE Configuration

IN THIS SECTION

- Requirements | 288
- Overview | 288
- Configuration | 288
- Verification | 292

Generic routing encapsulation (GRE) is an IP encapsulation protocol that is used to transport packets over a network. Information is sent from one network to the other through a GRE tunnel. GRE encapsulates a payload as a GRE packet. This GRE packet is encapsulated in an outer protocol (delivery protocol). GRE tunnel endpoints forward payloads into GRE tunnels for routing packets to the destination. After reaching the end point, GRE encapsulation is removed and the payload is transmitted to its final destination. The primary use of GRE is to carry non-IP packets through an IP network; however, GRE is also used to carry IP packets through an IP cloud.

Requirements

- Configure a GRE (gr-) interface. The gr- interface contains a local address and destination address. It comes up as soon as it is configured. You can even configure an IP address on the gr- interface.
- Configure a route to reach the destination subnet (end-to-end connectivity). You can configure either a static route through the gr- interface or use an interior gateway protocol (IGP) such as OSPF.

Overview

GRE tunnels are designed to be completely stateless, which means that each tunnel endpoint does not keep any information about the state or availability of the remote tunnel endpoint. Normally, a GRE tunnel interface comes up as soon as it is configured, and it stays up as long as there is a valid tunnel source address or interface that is up.

Configuration

IN THIS SECTION

- Configuring a Route to Reach the Destination Subset | 289

By default, the local subnet interface is ge-0/0/0 with IPv4 address as 10.10.11.1/24. The destination subnet is 10.10.10.0/24 with the tunnel endpoint IPv4 interface as 10.10.10.1/24.

GRE configuration shows the default configuration between the tunnel interfaces on SRX series devices.

Figure 18: GRE Configuration

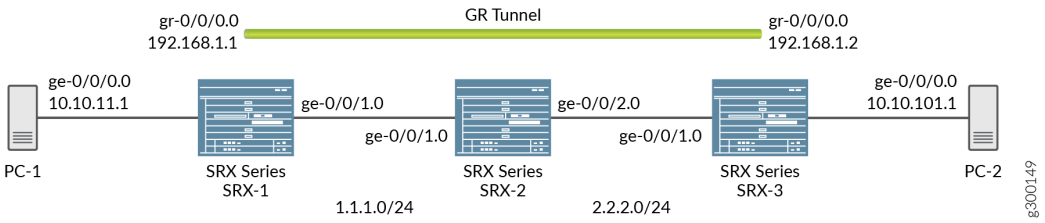

Configuring a Route to Reach the Destination Subset

Step-by-Step Procedure

You can either configure a static route through the gr- interface or by using IGP.

1. Configure the local subnet interface ge-0/0/0 interface.

```
[edit interfaces]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.10.11.1/24
```

2. Configure the interface ge-0/0/1.

```
[edit interfaces]
user@host# set interfaces ge-0/0/1 unit 0 family inet address 1.1.1.1/24
```

3. Configure the gr- tunnel endpoints and specify the source address, destination address, and family as inet for the tunnel endpoints.

```
[edit interfaces]
user@host# set interfaces gr-0/0/0  unit 0 tunnel source 1.1.1.1 destination 2.2.2.1
user@host# set interfaces gr-0/0/0 unit 0 family inet address 192.168.1.1/24
```

4. The configured interfaces are bound to a security zone at the `[edit security]` hierarchy level. Use the `show zones` command to view the zones. Configure the zones as follows:

```
[edit security zones security-zones trust]]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces gr-0/0/0.0
user@host# set zones zone names protocols all
```

```
[edit security zones security-zones untrust]]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
```

5. View the configured interfaces at the `[edit interfaces]` hierarchy level using the `show` command.

```
[edit interfaces]
user@host# set routing options static route 10.10.10.0/24 next hop gr-0/0/0.0
```

6. In case you do not want to define a static route, OSPF can be configured between gr-0/0/0 interfaces on both the sides and internal subnet as passive neighbor, to receive all the internal routes. Configure OSPF at the `[edit protocols]` hierarchy level and view it using the `show` command.

```
[edit protocols]
user@host# set protocols ospf area 0.0.0.0 interface gr-0/0/0.0
```

Results

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

GRE configuration using the static route:

```
[edit interfaces]
root@SRX-1# show
ge-0/0/0 {
    unit 0 {
        family inet {
```

```
            address 10.10.11.1/24;
        }
    }
}

gr-0/0/0 {
    unit 0 {
        tunnel {
            source 1.1.1.1;
            destination 2.2.2.1;
        }
        family inet {
            address 192.168.1.1/24;
        }
    }
}

ge-0/0/1 {
    unit 0 {
        family inet {
            address 1.1.1.1/24;
        }
    }
}

[edit security]
root@SRX-1# show
zones {
    security-zone trust {
        host-inbound-traffic {
            system-services {
                all;
            }
            protocols {
                all;
            }
        }
        interfaces {
            gr-0/0/0.0;
        }
    }
}
```

```
root@SRX-1# show routing-options
static {
    route 10.10.10.0/24 next-hop gr-0/0/0.0;
}
```

GRE configuration using OSPF configured between interfaces gr-0/0/0 on both sides and internal subnet as passive neighbor:

```
[edit protocols]
root@SRX-1# show
ospf {
    area 0.0.0.0 {
        interface gr-0/0/0.0;
        interface ge-0/0/0.0 {
            passive;
        }
    }
}
```

Verification

IN THIS SECTION

- Verification of the GRE Interfaces | **292**
- Verification of the Route | **293**
- Verification of Traffic Through GRE Tunnel | **293**

To verify that the configuration of GRE on the SRX Series device is successful, perform the following tasks:

Verification of the GRE Interfaces

Purpose

Verify that the GRE interfaces are up.

Action

Run the `show interfaces` command at the `[edit interfaces]` hierarchy level:

```
show interfaces gr-0/0/0 terse
[edit interfaces]
Interface Admin Link Proto Local Remote
gr-0/0/0   up up
gr-0/0/0.0 up up inet 192.168.1.1/24
```

Verification of the Route

Purpose

Verify that the route for the destination network is reachable through the GRE tunnel interface.

Action

Run the `show route forwarding-table matching` *`10.10.10.0/24`* command at the `[edit interfaces]` hierarchy level:

```
[edit interfaces]
user@router# run show route forwarding-table matching 10.10.10.0/24
Routing table: default.inet
Internet:
 ....
Destination        Type RtRef Next hop           Type Index    NhRef Netif
10.10.10.0/24      user     0                    ucst      595     2 gr-0/0/0.0
```

Verification of Traffic Through GRE Tunnel

Purpose

Send the traffic to the destination subnet and verify when the GRE interface is up.

Action

Run the `show interfaces gr-0/0/0 extensive` operational command. Also verify that the packets are leaving through the gr- interface.

```
user@host> show interfaces gr-0/0/0 extensive
Physical interface: gr-0/0/0, Enabled, Physical link is Up
Interface index: 134, SNMP ifIndex: 40, Generation: 17
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: 2005-08-05 21:39:41 UTC (00:00:47 ago)
Traffic statistics:
Input bytes : 8400 0 bps
Output bytes : 8400 0 bps
Input packets: 100 0 pps
Output packets: 100 0 pps

Logical interface gr-0/0/0.0 (Index 72) (SNMP ifIndex 28) (Generation 17)
Flags: Point-To-Point SNMP-Traps 16384
IP-Header 10.1.1.2:10.1.1.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Traffic statistics:
Input bytes : 8400
Output bytes : 8400
Input packets: 100
Output packets: 100
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 8400 0 bps
Output bytes : 8400 0 bps
Input packets: 100 0 pps
Output packets: 100 0 pps
Protocol inet, MTU: 1476, Generation: 25, Route table: 0
Flags: None
Addresses, Flags: Is-Primary
```

```
Destination: Unspecified, Local: 100.1.1.1, Broadcast: Unspecified,
Generation: 30
```

SEE ALSO

Generic Routing Encapsulation (GRE)

Understanding Generic Routing Encapsulation

Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

Example: Configuring GRE over IPsec Tunnels

IN THIS SECTION

- Requirements | **295**
- Overview | **295**
- Configuration | **296**
- Verification | **299**

Requirements

Overview

GRE tunnels offer minimal security, whereas an IPsec tunnel offers enhanced security in terms of confidentiality, data authentication, and integrity assurance. Also, IPsec cannot directly support multicast packets. However, if an encapsulated GRE tunnel is used first, an IPsec tunnel can then be used to provide security to the multicast packet. In a GRE over IPsec tunnel, all of the routing traffic (IP and non-IP) can be routed through. When the original packet (IP/non-IP) is GRE encapsulated, it has an IP header as defined by the GRE tunnel, normally the tunnel interface IP addresses. The IPsec protocol can understand the IP packet; so it encapsulates the GRE packet to make it GRE over IPsec.

The basic steps involved in configuring GRE over IPsec are as follows:

- Configure the route-based IPsec tunnel.
- Configure the GRE tunnel.

- Configure a static route with the destination as the remote subnet through the gr- interface.
- Configure the static route for the GRE endpoint with the st0 interface as next hop.

Configuration

IN THIS SECTION

- Configuring a GRE interface over an IPsec tunnel | **296**
- Results | **297**

In this example, the default configuration has the local subnet interface as ge-0/0/0 with the IPv4 address as 10.10.11.1/24. The destination subnet is 10.10.10.0/24. The gr-0/0/0 interface tunnel endpoints are loopback addresses on both the sides, with the local loopback IPv4 address as 172.20.1.1 and the remote loopback IPv4 address as 172.20.1.2. The gr-0/0/0, st0 and lo0 interfaces are bound to a security zone and policies are created accordingly.

Configuring a GRE interface over an IPsec tunnel

Step-by-Step Procedure

1. Configure the GRE at the [set interfaces *interface-name* unit *unit-number*] hierarchy level, where the interface name is ge-0/0/0, and the family is set as inet.

```
[edit interfaces]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.10.11.1/24
```

2. Configure the gr- tunnel endpoints and specify the source address, destination address, and family as inet for the tunnel endpoints.

```
[edit interfaces]
user@host# set interfaces gr-0/0/0  unit 0 tunnel source 172.20.1.1 destination 172.20.1.2
user@host# set interfaces gr-0/0/0  unit 0 family inet 192.168.1.1/24
```

3. Similarly configure the lo0 and st0 interface with the family set as inet.

```
[edit interfaces]
user@host# set interfaces lo0 unit 0 family inet address 172.20.1.1/32
```

```
[edit interfaces]
user@host# set interfaces st0 unit 0 family inet
```

4. Confifure the GRE interfaces with security zones. Use the `show zones` command to view the zones, where the configured tunnel interfaces, lo0 and st0 are displayed.

```
[edit security zones security-zones trust]]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces gr-0/0/0.0
user@host# set zones zone names protocols all
user@host# set interfaces lo0.0
user@host# set interfaces st0.0
```

```
[edit security zones security-zones untrust]]
user@host# set host-inbound-traffic system-services all
user@host# set host-inbound-traffic protocols all
user@host# set interfaces gr-0/0/0.0.1
user@host# set interfaces lo0.0
user@host# set interfaces st0.0
```

Results

In configuration mode, confirm your interface configuration by entering the `show` command. The configured interfaces are bound to a security zone at the `[edit security]` hierarchy level. Use the `show zones` command to view the zones, where the configured interfaces (gr-, st0.0, and lo0) are displayed. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Parameters for configuring the GRE interfaces:

```
user@host> show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.10.11.1/24;
        }
    }
}

gr-0/0/0 {
    unit 0 {
        tunnel {
            source 172.20.1.1;
            destination 172.20.1.2;
        }
        family inet {
            address 192.168.1.1/24;
        }
    }
}

lo0 {
    unit 0 {
        family inet {
            address 172.20.1.1/32;
        }
    }
}

st0 {
    unit 0 {
        family inet;
    }
}

[edit]
root@Juniper# show
routing-options {
    static {
        route 10.10.10.0/24 next-hop gr-0/0/0.0;
```

```
        route 172.20.1.2/32 next-hop st0.0;
    }
}
```

Parameters for configuring the GRE interfaces with security zones:

```
[edit security]
root@Juniper# show
zones {
    security-zone trust {
        host-inbound-traffic {
            system-services {
                all;
            }
            protocols {
                all;
            }
        }
        interfaces {
            gr-0/0/0.0;
            lo0.0;
            st0.0;
        }
    }
}
```

Verification

IN THIS SECTION

- Verification of the IPsec Tunnel | **299**

Verification of the IPsec Tunnel

Purpose

Verify that the IPsec tunnel is up.

Action

Run the commands `show security ike security-associations` and `show security ipsec security-associations` commands.

SEE ALSO

Generic Routing Encapsulation (GRE)

Understanding Generic Routing Encapsulation

Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

Example: Configuring a GRE Tunnel When the Tunnel Destination Is in a Routing Instance

IN THIS SECTION

- Requirements | **300**
- Overview | **300**
- Configuration | **301**
- Verification | **306**

Requirements

Overview

You can configure a GRE tunnel when the tunnel destination is in a default routing instance or non-default routing instance. Configuration of a GRE tunnel requires defining the tunnel source and the tunnel destination addresses. If the tunnel destination is in a routing instance, and there is more than one routing instance present, you need to specify the correct routing instance and also the routing table to be used to reach the configured tunnel destination address.

NOTE: The tunnel destination address is by default considered to be reachable using the default routing table "inet.0".

Configuration

IN THIS SECTION

- Configuring a GRE Tunnel When the Tunnel Destination Is in a Default Routing Instance | **301**
- Configuring a GRE Tunnel When the Tunnel Destination Is in a Non-default Routing Instance | **302**
- Results | **303**

In this example, you can configure a GRE tunnel between the gr- interfaces on SRX Series devices with two instances. The instances are when the tunnel destination is in a default routing instance and when the tunnel destination is in a non-default routing instance.

Configuring a GRE Tunnel When the Tunnel Destination Is in a Default Routing Instance

This example uses the default routing instance to reach the tunnel destination. Because of this, the routing table inet.0 is used by default.

Step-by-Step Procedure

1. Specify the source and destination address of the tunnel.

```
[edit interfaces]
user@host# set interfaces gr-0/0/0  unit 0 tunnel source 172.16.0.1 destination 10.10.1.2
user@host# set interfaces gr-0/0/0  unit 0 family inet 192.168.100.1/30;
```

2. Configure the ge- interface and lo0 interface with the family set as inet.

```
[edit interfaces]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 172.30.73.56/24
user@host# set interfaces lo0 unit 0 family inet address 172.16.0.1/32
```

3. Configure the GRE tunnel interface for routing options as mentioned in the GRE configuration topic.

Configuring a GRE Tunnel When the Tunnel Destination Is in a Non-default Routing Instance

For a non-default routing instance, ensure that you have already configured the gr-0/0/0 interface.

Step-by-Step Procedure

1. Configure the GRE tunnel with the gr-0/00 interface and family set as inet.

```
[edit interfaces]
user@host# set interfaces gr-0/00 unit 0 family inet address
```

2. Specify the source and destination address of the tunnel.

```
[edit interfaces]
user@host# set interfaces gr-0/0/0  unit 0 tunnel source 172.16.0.1 tunnel destination
10.10.1.2 family inet 192.168.100.1/30;
```

3. Configure the ge- interface and lo0 interface with the family set as inet.

```
[edit interfaces]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 172.30.73.56/24
user@host# set interfaces lo0 unit 0 family inet address 172.16.0.1/32
```

4. Configure the routing instances used for the tunnel interface.

```
[edit routing-instances]
user@host# set routing-instances test instance-type virtual-router
user@host# set routing-instances test routing-options static route 10.10.1.2/32 next-hop
172.30.73.57
user@host# set routing-instances test interface ge-0/0/0.0
```

5. Configure the routing-instance for GRE tunnel interfaces.

```
[edit interfaces]
user@host# set interfaces gr-0/0/0 unit 0  tunnel routing-instance destination test
```

6. Add the static route for tunnel destination.

```
[edit interfaces]
user@host# set routing-options static route 10.10.1.2/32 next-table test.inet.0
```

NOTE: When the SRX Series device is in packet mode, you do not need to configure a static route to make the tunnel destination reachable from inet.0. However, you still need to specify the correct routing instance under the gr-0/0/0 interface.

Results

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

When the tunnel destination is in a default routing instance:

```
interfaces {
    gr-0/0/0 {
        unit 0 {
            tunnel {
                source 172.16.0.1;
                destination 10.10.1.2;
            }
            family inet {
                address 192.168.100.1/30;
            }
        }
    }
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.30.73.56/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
```

```
                address 172.16.0.1/32;
            }
        }
    }
    ...
}
routing-options {
    static {
        route 10.10.1.2/32 next-hop 172.30.73.57;                    # Tunnel destination is
reachable from default routing-instance
        ...
    }
}
routing-instances {
    test {
        instance-type virtual-router;
        interface gr-0/0/0.0;
        routing-options {
            ...
        }
    }
}
```

When the tunnel destination is in a non-default routing instance:

```
interfaces {
    gr-0/0/0 {
        unit 0 {
            tunnel {
                source 172.16.0.1;
                destination 10.10.1.2;
                routing-instance {
                    destination test;                                # Routing-instance to reach
tunnel destination
                }
            }
            family inet {
                address 192.168.100.1/30;
            }
        }
    }
    ge-0/0/0 {
```

```
        unit 0 {
            family inet {
                address 172.30.73.56/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.0.1/32;
            }
        }
    }
    ...
}

routing-options {
         static {
         route 10.10.1.2/32 next-table test.inet.0;                # Tunnel destination is
reachable via test.inet.0
     ...
     }
}

routing-instances {
     test {
         instance-type virtual-router;
         interface ge-0/0/0;
         routing-options {
             static {
                 route 10.10.1.2/32 next-hop 172.30.73.57;         # Tunnel destination is
reachable from non-default routing-instance
                 ...
             }
         }
     }
}
```

Verification

IN THIS SECTION

- Verification of Static Route Use | 306
- Verification of Static Route Used in Default Instance | 307

Verification of Static Route Use

Purpose

Verify that the static route is used.

Action

Run the `show route forwarding table` command.

```
user@host> show route forwarding-table table test
No Title
Routing table: test.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    rjct      597     1
0.0.0.0/32         perm     0                    dscd      590     1
10.10.1.2/32       user     1 172.30.73.57       hold      598     4 ge-0/0/0.0
172.16.0.1.10.10.1.2.47/72
                   dest     0                    locl      617     1
172.30.73.0/24     intf     0                    rslv      588     1 ge-0/0/0.0
172.30.73.0/32     dest     0 172.30.73.0        recv      586     1 ge-0/0/0.0
172.30.73.56/32    intf     0 172.30.73.56       locl      587     2
172.30.73.56/32    dest     0 172.30.73.56       locl      587     2
172.30.73.57/32    dest     0 172.30.73.57       hold      598     4 ge-0/0/0.0
172.30.73.255/32   dest     0 172.30.73.255      bcst      585     1 ge-0/0/0.0
224.0.0.0/4        perm     0                    mdsc      596     1
224.0.0.1/32       perm     0 224.0.0.1          mcst      600     1
255.255.255.255/32 perm     0                    bcst      601     1
```

Verification of Static Route Used in Default Instance

Purpose

Verify that the static route is used for the default instance.

Action

Run the `show route forwarding table` command.

```
user@host> show route forwarding-table matching 10.10.1.2
Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
10.10.1.2/32       user     0                    rtbl      604     3
```

SEE ALSO

- *Generic Routing Encapsulation (GRE)*
- *Understanding Generic Routing Encapsulation*
- *Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly*

RELATED DOCUMENTATION

- *Generic Routing Encapsulation (GRE)*

Configuring Point-to-Point Protocol over Ethernet

IN THIS SECTION

- Understanding Point-to-Point Protocol over Ethernet | 308
- Understanding PPPoE Interfaces | 312

- Example: Configuring PPPoE Interfaces | **312**
- Understanding PPPoE Ethernet Interfaces | **322**
- Example: Configuring PPPoE Encapsulation on an Ethernet Interface | **322**
- Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces | **324**
- Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface | **324**
- Understanding CHAP Authentication on a PPPoE Interface | **328**
- Example: Configuring CHAP Authentication on a PPPoE Interface | **328**
- Verifying Credit-Flow Control | **331**
- Verifying PPPoE Interfaces | **333**
- Verifying R2CP Interfaces | **334**
- Displaying Statistics for PPPoE | **336**
- Setting Tracing Options for PPPoE | **337**

Point-to-Point Protocol over Ethernet (PPPoE) combines PPP, with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator. The below topics discuss the overview of PPPoE interfaces, PPPoE Ethernet interfaces, PPPoE ATM-over-ADSL, and ATM-over-SHDSL Interfaces, CHAP aunthentication on PPPoE, displaying statistics, setting tracing options for PPPoE and verification of these interfaces on security devices.

Understanding Point-to-Point Protocol over Ethernet

IN THIS SECTION

- PPPoE Discovery Stage | **309**
- PPPoE Session Stage | **310**

Point-to-Point Protocol over Ethernet (PPPoE) combines PPP, which typically runs over broadband connections, with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator. PPPoE enables service providers to maintain access control through PPP connections and also manage multiple hosts at a remote site.

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

To use PPPoE, you must initiate a PPPoE session, encapsulate Point-to-Point Protocol (PPP) packets over Ethernet, and configure the device as a PPPoE client. To provide a PPPoE connection, each PPP session must learn the Ethernet address of the remote peer and establish a unique session identifier during the PPPoE discovery and session stages.

NOTE: Juniper Networks devices with asymmetric digital subscriber line (ADSL) or symmetric high-speed DSL (SHDSL) interfaces can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections.

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

This topic contains the following sections:

PPPoE Discovery Stage

To initiate a PPPoE session, a host must first identify the Ethernet MAC address of the remote peer and establish a unique PPPoE session ID for the session. Learning the remote Ethernet MAC address is called PPPoE discovery.

During the PPPoE discovery process, the host does not discover a remote endpoint on the Ethernet network. Instead, the host discovers the access concentrator through which all PPPoE sessions are established. Discovery is a client/server relationship, with the host (a device running Junos OS) acting as the client and the access concentrator acting as the server. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.

NOTE: A device cannot receive PPPoE packets from two different access concentrators on the same physical interface.

The PPPoE discovery stage consists of the following steps:

1. PPPoE Active Discovery Initiation (PADI)—The client initiates a session by broadcasting a PADI packet to the LAN to request a service.

2. PPPoE Active Discovery Offer (PADO)—Any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client.

3. PPPoE Active Discovery Request (PADR)—From the PADOs it receives, the client selects one access concentrator based on its name or the services offered and sends it a PADR packet to indicate the service or services needed.

4. PPPoE Active Discovery Session-Confirmation (PADS)—When the selected access concentrator receives the PADR packet, it accepts or rejects the PPPoE session:

   - To accept the session, the access concentrator sends the client a PADS packet with a unique session ID for a PPPoE session and a service name that identifies the service under which it accepts the session.

   - To reject the session, the access concentrator sends the client a PADS packet with a service name error and resets the session ID to zero.

PPPoE Session Stage

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

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

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

NOTE: If PPPoE session is already up and the user restarts the PPPoE daemon, a new PPPoE daemon with a new PID starts while the existing session is not terminated.

If PPPoE session is already down and user restarts the PPPoE daemon, the PPPoE discovery establishes a new session.

The PPPoE session is not terminated for the following configuration changes:

- Changing idle time out value
- Changing auto rec timer value
- Deleting idle time out
- Deleting auto rec timer
- Add new auto rec time
- Add new idle time out
- Change negotiate address to static address
- Change static ip address to a new static ip address
- Changing default chap secrete

The PPPoE session is terminated for the following configuration changes:

- Add ac name
- Delete chap ppp options
- Add new chap ppp options
- Configure uifd mac

> **NOTE**: When the MTU for an underlying physical interface is changed, it brings down the PPPoE session. The PPPoE MTU can be greater than 1492 if the Ethernet or WAN connection supports RFC 4638 (Mini Jumbo Frames).

SEE ALSO

Understanding Physical Encapsulation on an Interface | 274

Understanding PPPoE Interfaces | 312

Understanding PPPoE Ethernet Interfaces | 322

Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces | 324

Understanding CHAP Authentication on a PPPoE Interface | 328

Understanding the PPPoE-Based Radio-to-Router Protocol

Understanding PPPoE Interfaces

The device's Point-to-Point Protocol over Ethernet (PPPoE) interface to the access concentrator can be a Fast Ethernet interface, a Gigabit Ethernet interface, a redundant Ethernet interface, an ATM-over-ADSL interface, or an ATM-over-SHDSL interface. The PPPoE configuration is the same for all interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:

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

To configure a PPPoE interface, you create an interface with a *logical interface* unit 0, then specify a logical Ethernet or ATM interface as the underlying interface for the PPPoE session. You then specify other PPPoE options, including the access concentrator and PPPoE session parameters.

NOTE: PPPoE over redundant Ethernet (reth) interface is supported on SRX100, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340 and SRX650 devices. (Platform support depends on the Junos OS release in your installation.) This feature allows an existing PPPoE session to continue without starting a new PPP0E session in the event of a failover.

Example: Configuring PPPoE Interfaces

IN THIS SECTION

- Requirements | **312**
- Overview | **313**
- Configuration | **313**
- Disabling the End-of-List Tag | **319**

This example shows how to configure a PPPoE interface.

Requirements

Before you begin, configure an Ethernet interface. See "Example: Creating an Ethernet Interface" on page 166.

Overview

In this example, you create the PPPoE interface pp0.0 and specify the logical Ethernet interface ge-0/0/1.0 as the underlying interface. You also set the access concentrator, set the PPPoE session parameters, and set the MTU of the IPv4 family to 1492.

Configuration

IN THIS SECTION

- Verification | 315

Procedure

CLI Quick Configuration

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

```
set interfaces pp0 unit 0 pppoe-options underlying-interface ge-0/0/1.0 access-concentrator
ispl.com auto-reconnect 100 idle-timeout 100 client service-name video@ispl.com
set interfaces pp0 unit 0  family inet mtu 1492 negotiate-address
```

Step-by-Step Procedure

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

To configure a PPPoE interface:

1. Create a PPPoE interface.

```
[edit]
user@host# edit interfaces pp0 unit 0
```

2. Configure PPPoE options.

```
[edit interfaces pp0 unit 0]
user@host# set pppoe-options underlying-interface ge-0/0/1.0 access-concentrator ispl.com
auto-reconnect 100 idle-timeout 100 client service-name video@ispl.com
```

3. Configure the MTU.

```
[edit interfaces pp0 unit 0]
user@host# set family inet mtu 1492
```

NOTE: If you want to configure `mtu` to a value above 1492 octets, then use `ppp-max-payload` option. Refer *pppoe-options* for more details.

4. Configure the PPPoE interface address.

```
[edit interfaces pp0 unit 0]
user@host# set family inet negotiate-address
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces pp0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces pp0
unit 0 {
    pppoe-options {
        underlying-interface ge-0/0/1.0;
        idle-timeout 100;
        access-concentrator ispl.com;
        service-name "vide0@ispl.com";
        auto-reconnect 100;
        client;
     }
    family inet {
```

```
        mtu 1492;
        negotiate-address;
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying PPPoE Interfaces | **315**
- Verifying PPPoE Sessions | **317**
- Verifying the PPPoE Version | **317**
- Verifying PPPoE Statistics | **318**

Confirm that the configuration is working properly.

Verifying PPPoE Interfaces

Purpose

Verify that the PPPoE device interfaces are configured properly.

Action

From operational mode, enter the `show interfaces pp0` command.

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 67, SNMP ifIndex: 317
  Type: PPPoE, Link-level type: PPPoE, MTU: 9192
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
```

```
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)

Logical interface pp0.0 (Index 1) (SNMP ifIndex 330)
  Flags: Point-To-Point SNMP-Traps 16384 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 3304,
    Session AC name: isp1.com, AC MAC address: 00:90:1a:40:f6:4c,
    Service name: video@isp1.com, Configured AC name: isp1.com,
    Auto-reconnect timeout: 60 seconds
    Underlying interface: ge-5/0/0.0 (Index 71)
Input packets : 23
Output packets: 22
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 16 (00:00:26 ago), Output: 0 (never)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Success
  Protocol inet, MTU: 1492
    Flags: Negotiate-Address
    Addresses, Flags: Kernel Is-Preferred Is-Primary
      Destination: 211.211.211.2, Local: 211.211.211.1
```

The output shows information about the physical and the logical interfaces. Verify the following information:

- The physical interface is enabled and the link is up.
- The PPPoE session is running on the correct logical interface.
- For state, the state is active (up).
- For underlying interface, the physical interface on which the PPPoE session is running is correct:
  - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, ge-5/0/0.0.
  - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, at-2/0/0.0.

Verifying PPPoE Sessions

Purpose

Verify that a PPPoE session is running properly on the logical interface.

Action

From operational mode, enter the `show pppoe interfaces` command.

```
user@host> show pppoe interfaces
pp0.0 Index 67
  State: Session up, Session ID: 31,
  Service name: video@isp1.com, Configured AC name: isp1.com,
  Session AC name: belur, AC MAC address: 00:90:1a:40:f6:4e,
  Auto-reconnect timeout: 1 seconds,
  Underlying interface: ge-0/0/1.0 Index 69
```

The output shows information about the PPPoE sessions. Verify the following information:

- The PPPoE session is running on the correct logical interface.
- For state, the session is active (up).
- For underlying interface, the physical interface on which the PPPoE session is running is correct:
  - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, ge-0/0/1.0.
  - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, at-2/0/0.0.

NOTE: To clear a PPPoE session on the pp0.0 interface, use the `clear pppoe sessions pp0.0` command. To clear all sessions on the interface, use the `clear pppoe sessions` command.

Verifying the PPPoE Version

Purpose

Verify the version information of the PPPoE protocol configured on the device interfaces.

Action

From operational mode, enter the `show pppoe version` command.

```
user@host> show pppoe version
Point-to-Point Protocol Over Ethernet, version 1. rfc2516
   PPPoE protocol               = Enabled
   Maximum Sessions             = 256
   PADI resend timeout          = 2 seconds
   PADR resend timeout          = 16 seconds
   Max resend timeout           = 64 seconds
   Max Configured AC timeout    = 4 seconds
```

The output shows PPPoE protocol information. Verify the following information:

- The correct version of the PPPoE protocol is configured on the interface.
- For PPPoE protocol, the PPPoE protocol is enabled.

Verifying PPPoE Statistics

Purpose

Verify the statistics information about PPPoE interfaces.

Action

From operational mode, enter the `show pppoe statistics` command.

```
user@host> show pppoe statistics
Active PPPoE sessions: 4
  PacketType                      Sent         Received
    PADI                           502                0
    PADO                             0              219
    PADR                           219                0
    PADS                             0              219
    PADT                             0              161
    Service name error               0                0
    AC system error                  0               13
    Generic error                    0                0
    Malformed packets                0               41
```

```
    Unknown packets                  0                0
  Timeout
    PADI                            42
    PADO                             0
    PADR                             0
```

The output shows information about active sessions on PPPoE interfaces. Verify the following information:

- Total number of active PPPoE sessions running on the interfac
- For packet type, the number of packets of each type sent and received during the PPPoE session

Disabling the End-of-List Tag

IN THIS SECTION

- Procedure | 319
- Verifying That the End-of-List Tag Is Disabled | 320

During the PPPoE discovery stage, any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client. When a client receives a PADO packet, and if it encounters the `End-of-List` tag in the PADO packet, tags after the `End-of-List` tag are ignored and the complete information is not processed correctly. As a result, the PPPoE connection is not established correctly.

Starting in Junos OS Release 12.3X48-D10 you can avoid some PPPoE connection errors by configuring the `ignore-eol-tag` option to disable the `End-of-List` tag in the PADO packet.

Procedure

Step-by-Step Procedure

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

To disable the `End-of-List` tag:

1. Create a PPPoE interface.

```
[edit]
user@host# set interfaces pp0 unit 0
```

2. Configure PPPoE options.

```
[edit interfaces pp0 unit 0]
user@host# set pppoe-options ignore-eol-tag
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces pp0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces pp0
unit 0 {
pppoe-options {
    ignore-eol-tag;
 }
```

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

Verifying That the End-of-List Tag Is Disabled

Purpose

Verify the status of the `End-of-List` tag in the PPPoE configuration.

Action

From operational mode, enter the `show interfaces pp0.0` command.

```
user@host> show pppoe interfaces pp0.0
Logical interface pp0.0 (Index 78) (SNMP ifIndex 541)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
```

```
    PPPoE:
      State: SessionUp, Session ID: 3,
      Session AC name: cell, Remote MAC address: 00:26:88:f7:77:83,
      Configured AC name: None, Service name: None,
      Auto-reconnect timeout: Never, Idle timeout: Never,
      Underlying interface: ge-0/0/3.0 (Index 77)
     Ignore End-Of-List tag: Enable
```

```
user@host> show pppoe interfaces pp0.0 extensive
pp0.0 Index 74
  State: Session up, Session ID: 1,
  Service name: None,
  Session AC name: cell, Configured AC name: None,
  Remote MAC address: 00:26:88:f7:77:83,
  Session uptime: 00:02:03 ago,
  Auto-reconnect timeout: 10 seconds, Idle timeout: Never,
  Underlying interface: ge-0/0/3.0 Index 73
  Ignore End-of-List tag: Enable
  PacketType                      Sent         Received
    PADI                            23                0
    PADO                             0                5
    PADR                            11                0
    PADS                             0                2
    PADT                             2                0
    Service name error               0                0
    AC system error                  0                0
    Generic error                    0                0
    Malformed packets                0                0
    Unknown packets                  0                0
  Timeout
    PADI                             3
    PADO                             0
    PADR                             3
  Receive Error Counters
    PADI                             0
    PADO                             0
    PADR                             0
    PADS                             0
```

The output shows information about active sessions on PPPoE interfaces. Verify that the `Ignore End-of-List tag: Enable` option is set.

Understanding PPPoE Ethernet Interfaces

During a Point-to-Point Protocol over Ethernet (PPPoE) session, the device encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 19 on page 322 shows a typical PPPoE session between a device and an access concentrator on the Ethernet loop.

Figure 19: PPPoE Session on the Ethernet Loop

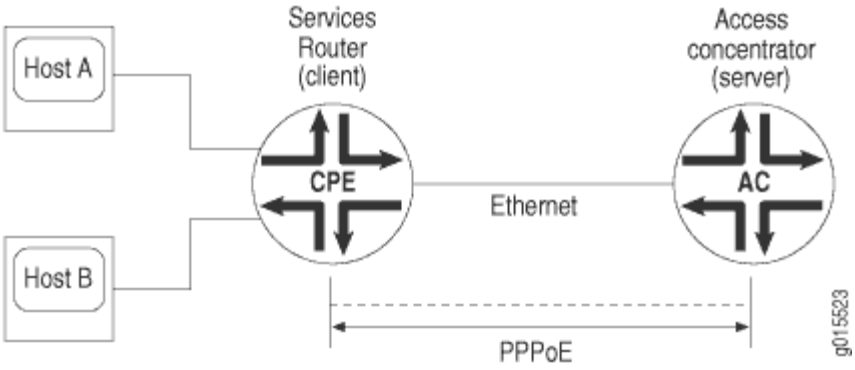

To configure PPPoE on an Ethernet interface, you configure encapsulation on the *logical interface*.

Example: Configuring PPPoE Encapsulation on an Ethernet Interface

IN THIS SECTION

- Requirements | 322
- Overview | 323
- Configuration | 323
- Verification | 323

This example shows how to configure PPPoE encapsulation on an Ethernet interface.

Requirements

Before you begin:

- Configure an Ethernet interface. See "Example: Creating an Ethernet Interface" on page 166.

- Configure a PPPoE encapsulation interface. See "Example: Configuring PPPoE Interfaces" on page 312.

Overview

In this example, you configure PPPoE encapsulation on the ge-0/0/1 interface.

Configuration

IN THIS SECTION

- Procedure | 323

Procedure

Step-by-Step Procedure

To configure PPPoE encapsulation:

1. Enable PPPoE encapsulation on the interface.

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 encapsulation ppp-over-ether
```

2. Commit the configuration if you are done configuring the device.

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

Verification

To verify the configuration is working properly, enter the `show interfaces ge-0/0/1` command.

Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces

When an ATM network is configured with a point-to-point connection, Point-to-Point Protocol over Ethernet (PPPoE) can use ATM Adaptation Layer 5 (AAL5) for framing PPPoE-encapsulated packets. The AAL5 protocol provides a virtual connection between the client and the server within the same network. The device encapsulates each PPPoE frame in an ATM frame and transports each frame over an asymmetric digital subscriber line (ADSL) or symmetric high-speed DSL (SHDSL) loop and a digital subscriber line access multiplexer (DSLAM). For example, Figure 20 on page 324 shows a typical PPPoE over ATM session between a device and an access concentrator on an ADSL loop.

Figure 20: PPPoE Session on an ADSL Loop

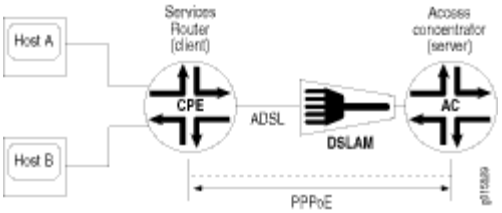

For PPPoE on an ATM-over-ADSL or ATM-over-SHDSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL or ATM-over-SHDSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL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

Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface

IN THIS SECTION

- Requirements | 325
- Overview | 325
- Configuration | 325
- Verification | 327

This example shows how to configure a physical interface for Ethernet over ATM encapsulation and how to create a logical interface for PPPoE over LLC encapsulation.

Requirements

Before you begin:

- Configure network interfaces. See "Example: Creating an Ethernet Interface" on page 166.
- Configure PPPoE interfaces. See "Example: Configuring PPPoE Interfaces" on page 312.
- Configure PPPoE encapsulation on an Ethernet interface. See "Example: Configuring PPPoE Encapsulation on an Ethernet Interface" on page 322.

Overview

In this example, you configure the physical interface at-2/0/0 for Ethernet over ATM encapsulation. As part of the configuration, you set the virtual path identifier (VPI) on an ATM-over-ADSL physical interface to 0, you set the ADSL operating mode to auto, and you set the encapsulation type to ATM-over-ADSL. Then you create a logical interface for PPPoE over LLC encapsulation.

Configuration

IN THIS SECTION

- Procedure | 325

Procedure

CLI Quick Configuration

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

```
set interfaces at-2/0/0 atm-options vpi 0
set interfaces at-2/0/0 dsl-options operating-mode auto
set interfaces at-2/0/0 encapsulation ethernet-over-atm
set interfaces at-2/0/0 unit 0 encapsulation ppp-over-ether-over-atm-llc vci 0.120
```

Step-by-Step Procedure

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

To configure PPPoE encapsulation on an ATM-over-ADSL interface:

1. Configure the physical interface.

```
[edit]
user@host# edit interfaces at-2/0/0
```

2. Set the VPI on the interface.

```
[edit interfaces at-2/0/0]
user@host# set atm-options vpi 0
```

3. Configure the ADSL operating mode.

```
[edit interfaces at-2/0/0]
user@host# set dsl-options operating-mode auto
```

4. Configure PPPoE encapsulation.

```
[edit interfaces at-2/0/0]
user@host# set encapsulation ethernet-over-atm
```

5. Create a logical interface and configure LLC encapsulation.

```
[edit interfaces at-2/0/0]
user@host# set unit 0 encapsulation ppp-over-ether-over-atm-llc vci 0.120
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces at-2/0/0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces at-2/0/0 {
    encapsulation ethernet-over-atm;
        atm-options {
        vpi 0;
    }
    dsl-options {
        operating-mode auto;
     }
    unit 0 {
        encapsulation ppp-over-ether-over-atm-llc;
        vci 0.120;
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface | **327**

Confirm that the configuration is working properly.

Verifying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface

Purpose

Verify the PPPoE configuration for an ATM-over-ADSL or ATM-over-SHDSL interface.

Action

From operational mode, enter the `show interfaces` command.

Understanding CHAP Authentication on a PPPoE Interface

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

If you set the `passive` option to handle incoming CHAP packets only, the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not set the `passive` option, the interface always challenges its peer.

You can configure Remote Authentication Dial-In User Service (RADIUS) authentication of PPP sessions using CHAP. CHAP enables you to send RADIUS messages through a routing instance to customer RADIUS servers in a private network.

Example: Configuring CHAP Authentication on a PPPoE Interface

IN THIS SECTION

- Requirements | 328
- Overview | 329
- Configuration | 329
- Verification | 331

This example shows how to configure CHAP authentication on a PPPoE interface.

Requirements

Before you begin:

- Configure an Ethernet interface. See "Example: Creating an Ethernet Interface" on page 166.
- Configure a PPPoE interface. See "Example: Configuring PPPoE Interfaces" on page 312.

- Configure PPPoE encapsulation on an ATM-over-ADSL interface. See "Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface" on page 324.

Overview

In this example, you configure a CHAP access profile, and then apply it to the PPPoE interface pp0. You also configure the hostname to be used in CHAP challenge and response packets, and set the passive option for handling incoming CHAP packets.

Configuration

IN THIS SECTION

- Procedure | 329

Procedure

CLI Quick Configuration

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

```
set access profile A-ppp-client client client1 chap-secret my-secret
set interfaces pp0 unit 0 ppp-options chap access-profile A-ppp-client local-name A-ge-0/0/1.0
passive
```

Step-by-Step Procedure

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

To configure CHAP on a PPPoE interface:

1. Configure a CHAP access profile.

```
[edit]
user@host# set access profile A-ppp-client client client1 chap-secret my-secret
```

2. Enable CHAP options on the interface.

```
[edit]
user@host# edit interfaces pp0 unit 0 ppp-options chap
```

3. Configure the CHAP access profile on the interface.

```
[edit interfaces pp0 unit 0 ppp-options chap]
user@host# set access-profile A-ppp-client
```

4. Configure a hostname for the CHAP challenge and response packets.

```
[edit interfaces pp0 unit 0 ppp-options chap]
user@host# set local-name A-ge-0/0/1.0
```

5. Set the passive option to handle incoming CHAP packets only.

```
[edit interfaces pp0 unit 0 ppp-options chap]
user@host# set passive
```

Results

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

```
[edit]
    user@host# show interfaces
    pp0 {
    unit 0 {
     ppp-options {
    chap {
    access-profile A-ppp-client;
    local-name A-ge-0/0/1.0;
    passive;
    }
    }
```

```
    }
}
```

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

Verification

IN THIS SECTION

- Verifying CHAP Authentication | **331**

Confirm that the configuration is working properly.

Verifying CHAP Authentication

Purpose

Verify that CHAP is enabled on the interface.

Action

From operational mode, enter the `show interfaces` command.

Verifying Credit-Flow Control

IN THIS SECTION

- Purpose | **332**
- Action | **332**

Purpose

Display PPPoE credit-flow control information about credits on each side of the PPPoE session when credit processing is enabled on the interface.

Action

```
user@host> show pppoe interface detail
```

```
pp0.51 Index 73
  State: Session up, Session ID: 3,
  Service name: None,
  Configured AC name: None, Session AC name: None,
  Remote MAC address: 00:22:83:84:2e:81,
  Session uptime: 00:05:48 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/4.1 Index 72
  PADG Credits: Local: 12345, Remote: 6789, Scale factor: 128 bytes
  PADQ Current bandwidth: 750 Kbps, Maximum 1000 Kbps
    Quality: 85, Resources 65, Latency 100 msec.
  Dynamic bandwidth: 3 Kbps
pp0.1000 Index 71
  State: Down, Session ID: 1,
  Service name: None,
  Configured AC name: None, Session AC name: None,
  Remote MAC address: 00:00:00:00:00:00,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  PADG Credits: enabled
  Dynamic bandwidth: enabled
```

Verifying PPPoE Interfaces

IN THIS SECTION

- Purpose | 333
- Action | 333

Purpose

Display PPPoE interfaces information.

Action

- To display PPPoE interface information:

```
user@host> show pppoe interfaces pp0.51 detail
```

```
pp0.51 Index 75
  State: Session up, Session ID: 1,
  Service name: None,
  Configured AC name: None, Session AC name: None,
  Remote MAC address: 00:11:22:33:44:55,
  Session uptime: 00:04:18 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  PADQ Current bandwidth: 750 Kbps, Maximum 1000 Kbps
    Quality: 85, Resources 65, Latency 100 msec.
  Dynamic bandwidth: 3 Kbps
```

- To display PPPoE terse interface information:

```
user@host> show pppoe interfaces terse pp0.51
```

```
Interface    Admin Link Proto    Local        Remote
          pp0.51       up  up   inet   5.1.1.1     --> 5.1.1.2
                                  inet6    fe80::21f:12ff:fed2:2918/64
                                 feee::5:1:1:1/126
```

Verifying R2CP Interfaces

IN THIS SECTION

- Purpose | 334
- Action | 334

Purpose

Display R2CP interfaces information.

Action

- To display R2CP interface information:

```
root@host> show r2cp interfaces
```

```
Interface: ge-0/0/3.51
  Nodes: 0
```

- To display R2CP information:

```
root@host> show r2cp radio extensive
```

```
    Node Packet Type           Sent        Received          Errors
    MIM                           -               1               0
    ROM                           1               -               -
    Heartbeats                    0               0               0
    Node Term                     0               0               0
    Node Term Ack                 0               0               -

    Heartbeat Timeouts               0
    Node Term Timeouts               0


    Session Packet Type        Sent        Received          Errors
    Init                          -               1               0
    Init ACK                      1               -               -
    Update                        -               0               0
    Terminate                     0               0               0
    Terminate ACK                 0               0               0

    Terminate Timeouts               0
```

- To display R2CP session information:

```
root@host> show r2cp sessions extensive
```

```
Session: 1
  Destination MAC address 01:02:03:04:05:06
   Status: Established VLANs 201
   Virtual channel:  2
   Session Update:  last received: 3.268 seconds
     Current bandwidth: 22000 Kbps, Maximum 22000 Kbps
     Quality: 100, Resources 100, Latency 100 msec.
   Effective bandwidth: 952 Kbps, last change: 51.484 seconds
```

```
     Updates below threshold: 1


    Session Packet Type        Sent        Received          Errors
    Init                          -               1               0
    Init ACK                      1               -               -
    Update                        -               0               0
    Terminate                     0               0               0
    Terminate ACK                 0               0               0

    Terminate Timeouts               0
```

Displaying Statistics for PPPoE

IN THIS SECTION

- Purpose | 336
- Action | 336

Purpose

Display PPPoE statistics.

Action

```
user@host> show interfaces pp0.51 statistics
```

```
Logical interface pp0.51 (Index 75) (SNMP ifIndex 137)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
    PPPoE:
      State: SessionUp, Session ID: 1,
      Session AC name: None, Remote MAC address: 00:22:83:84:2f:03,
      Underlying interface: ge-0/0/4.1 (Index 74)
```

```
    Input packets : 20865
    Output packets: 284636
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 943 (00:00:06 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  PAP state: Closed
    Security: Zone: Null
    Protocol inet, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 5.1.1.2, Local: 5.1.1.1
    Protocol inet6, MTU: 1492
      Flags: None
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fed2:2918
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: feee::5:1:1:0/126, Local: feee::5:1:1:1
```

Setting Tracing Options for PPPoE

To trace the operations of the router's PPPoE process, include the traceoptions statement at the [edit protocols pppoe] hierarchy level:

```
[edit protocols pppoe]
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag flag;
    level severity-level;
    no-remote-trace;
}
```

To specify more than one tracing operation, include multiple `flag` statements.

You can specify the following flags in the `traceoptions` statement:

- `all`—All areas of code

- config—Configuration code
- events—Event code
- gres—Gres code
- init—Initialization code
- interface-db—Interface database code
- memory—Memory management code
- protocol—PPPoE protocol processing code
- rtsock—Routing socket code
- session-db—Session management code
- signal—Signal handling code
- state—State handling code
- timer—Timer code
- ui—User interface code

Release History Table

Release	Description
12.3X48-D10	Starting in Junos OS Release 12.3X48-D10 you can avoid some PPPoE connection errors by configuring the ignore-eol-tag option to disable the End-of-List tag in the PADO packet.

CHAPTER 6

Configuring Link Services Interfaces

Configuring Link Services Interfaces | 340

Configuring Link Fragmentation and Interleaving | 372

Configuring Class-of-Service on Link Services Interfaces | 376

Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles | 391

Configuring Multilink Frame Relay | 397

Configuring Compressed Real-Time Transport Protocol | 409

Configuring Link Services Interfaces

IN THIS SECTION

- Link Services Interfaces Overview | 340
- Link Services Configuration Overview | 348
- Verifying the Link Services Interface | 349
- Understanding the Internal Interface LSQ-0/0/0 Configuration | 354
- Example: Upgrading from ls-0/0/0 to lsq-0/0/0 for Multilink Services | 355
- Troubleshooting the Link Services Interface | 359

Juniper Networks devices support link services on the `lsq-0/0/0` link services queuing interface which includes multilink services like MLPP, MLFR and CRTP. The topics below discuss the overview of link services, configuration details and verification of the link services on SRX series devices.

Link Services Interfaces Overview

IN THIS SECTION

- Services Available on a Link Services Interface | 341
- Link Services Exceptions | 342
- Configuring Multiclass MLPPP | 343
- Queuing with LFI | 344
- Compressed Real-Time Transport Protocol Overview | 345
- Configuring Fragmentation by Forwarding Class | 346
- Configuring Link-Layer Overhead | 347

Link services include the multilink services Multilink Point-to-Point Protocol (MLPPP), Multilink Frame Relay (MLFR), and Compressed Real-Time Transport Protocol (CRTP). Juniper Networks devices support link services on the `lsq-0/0/0` link services queuing interface.

You configure the link services queuing interface (`lsq-0/0/0`) on a Juniper Networks device to support multilink services and CRTP.

The link services queuing interface on SRX Series devices consists of services provided by the following interfaces on the Juniper Networks M Series and T Series routing platforms: multilink services interface (`ml-`*`fpc/pic/port`*), link services interface (`ls-`*`fpc/pic/port`*), and link services intelligent queuing interface (`lsq-`*`fpc/pic/port`*). Although the multilink services, link services, and link services intelligent queuing (IQ) interfaces on M Series and T Series routing platforms are installed on Physical Interface Cards (PICs), the link services queuing interface on SRX Series devices is an internal interface only and is not associated with a physical medium or *Physical Interface Module* (PIM).

NOTE: (`ls-`*`fpc/pic/port`*) is not supported on SRX Series devices.

This section contains the following topics.

Services Available on a Link Services Interface

The link services interface is a *logical interface* available by default. Table 36 on page 341 summarizes the services available on the interface.

Table 36: Services Available on a Link Services Interface

Services	Purpose	More Information
Multilink bundles by means of MLPPP and MLFR encapsulation	Aggregates multiple constituent links into one larger logical bundle to provide additional bandwidth, load balancing, and redundancy. **NOTE**: Dynamic call admission control (DCAC) configurations are not supported on Link Services Interfaces.	• "Example: Configuring an MLPPP Bundle" on page 392 • "Example: Configuring Multilink Frame Relay FRF.15" on page 398 • "Example: Configuring Multilink Frame Relay FRF.16" on page 403
Link fragmentation and interleaving (LFI)	Reduces delay and *jitter* on links by breaking up large data packets and interleaving delay-sensitive voice packets with the resulting smaller packets.	"Understanding Link Fragmentation and Interleaving Configuration" on page 373

Table 36: Services Available on a Link Services Interface *(Continued)*

Services	Purpose	More Information
Compressed Real-Time Transport Protocol (CRTP)	Reduces the overhead caused by Real-Time Transport Protocol (RTP) on voice and video packets.	"Compressed Real-Time Transport Protocol Overview" on page 345
Class-of-service (CoS) classifiers, forwarding classes, schedulers and scheduler maps, and shaping rates	Provides a higher priority to delay-sensitive packets—by configuring CoS, such as the following: • Classifiers—To classify different types of traffic, such as voice, data, and network control packets. • Forwarding classes—To direct different types of traffic to different output queues. • Fragmentation map—To define mapping between forwarding class and multilink class, and forwarding class and fragment threshold. In forwarding class and multilink class mapping, drop timeout can be configured. • Schedulers and scheduler maps—To define properties for the output queues such as delay-buffer, transmission rate, and transmission priority. • Shaping rate—To define certain bandwidth usage by an interface.	• "Example: Configuring Interface Shaping Rates" on page 389 • "Configuring Fragmentation by Forwarding Class" on page 346

Link Services Exceptions

The link and multilink services implementation on SRX Series devices is similar to the implementation on the M Series and T Series routing platforms, with the following exceptions:

- Support for link and multilink services are on the `lsq-0/0/0` interface instead of the `ml-`*`fpc/pic/port`*, `lsq-`*`fpc/pic/port`*, and `ls-`*`fpc/pic/port`* interfaces.

- When LFI is enabled, fragmented packets are queued in a round-robin fashion on the constituent links to enable per-packet and per-fragment load balancing. See "Queuing with LFI" on page 344.
- Support for per-unit scheduling is on all types of constituent links (on all types of interfaces).
- Support for Compressed Real-Time Transport Protocol (CRTP) is for both MLPPP and PPP.

Configuring Multiclass MLPPP

For `lsq-0/0/0` on Juniper Networks device, with MLPPP encapsulation, you can configure multiclass MLPPP. If you do not configure multiclass MLPPP, fragments from different classes cannot be interleaved. All fragments for a single packet must be sent before the fragments from another packet are sent. Non-fragmented packets can be interleaved between fragments of another packet to reduce latency seen by non-fragmented packets. In effect, latency-sensitive traffic is encapsulated as regular PPP traffic, and bulk traffic is encapsulated as multilink traffic. This model works as long as there is a single class of latency-sensitive traffic, and there is no high-priority traffic that takes precedence over latency-sensitive traffic. This approach to LFI, used on the Link Services PIC, supports only two levels of traffic priority, which is not sufficient to carry the four-to-eight forwarding classes that are supported by M series and T series routing platforms.

Multiclass MLPPP makes it possible to have multiple classes of latency-sensitive traffic that are carried over a single multilink bundle with bulk traffic. In effect, multiclass MLPPP allows different classes of traffic to have different latency guarantees. With multiclass MLPPP, you can map each forwarding class into a separate multilink class, thus preserving priority and latency guarantees.

> **NOTE**: Configuring both LFI and multiclass MLPPP on the same bundle is not necessary, nor is it supported, because multiclass MLPPP represents a superset of functionality. When you configure multiclass MLPPP, LFI is automatically enabled.
>
> The Junos OS PPP implementation does not support the negotiation of address field compression and protocol field compression PPP NCP options, which means that the software always sends a full 4-byte PPP header.
>
> The Junos OS implementation of multiclass MLPPP does not support compression of common header bytes.

Multiclass MLPPP greatly simplifies packet ordering issues that occur when multiple links are used. Without multiclass MLPPP, all voice traffic belonging to a single flow is hashed to a single link to avoid packet ordering issues. With multiclass MLPPP, you can assign voice traffic to a high-priority class, and you can use multiple links.

To configure multiclass MLPPP on a link services IQ interface, you must specify how many multilink classes should be negotiated when a link joins the bundle, and you must specify the mapping of a forwarding class into an multiclass MLPPP class.

To specify how many multilink classes should be negotiated when a link joins the bundle, include the `multilink-max-classes` statement:

```
multilink-max-classes number;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`
- `[edit logical-routers logical-router-name interfaces interface-name unit logical-unit-number]`

The number of multilink classes can be 1 through 8. The number of multilink classes for each forwarding class must not exceed the number of multilink classes to be negotiated.

To specify the mapping of a forwarding class into a multiclass MLPPP class, include the `multilink-class` statement at the `[edit class-of-service fragmentation-maps forwarding-class class-name]` hierarchy level:

```
edit class-of-service fragmentation-maps forwarding-class class-namemultilink-class number
```

The multilink class index number can be 0 through 7. The `multilink-class` statement and the `no-fragmentation` statement are mutually exclusive.

To view the number of multilink classes negotiated, issue the `show interfaces lsq-0/0/0.logical-unit-number detail` command.

Queuing with LFI

LFI or non-LFI packets are placed into queues on constituent links based on the queues in which they arrive. No changes in the queue number occur while the fragmented, non-fragmented, or LFI packets are being queued.

For example, assume that Queue Q0 is configured with fragmentation threshold 128, Q1 is configured with no fragmentation, and Q2 is configured with fragmentation threshold 512. Q0 is receiving stream of traffic with packet size 512. Q1 is receiving voice traffic of 64 bytes, and Q2 is receiving stream of traffic with 128-byte packets. Next the stream on Q0 gets fragmented and queued up into Q0 of a constituent link. Also, all packets on Q2 are queued up on Q0 on constituent link. The stream on Q1 is considered to be LFI because no fragmentation is configured. All the packets from Q0 and Q2 are queued up on Q0 of constituent link. All the packets from Q1 are queued up on Q2 of constituent link.

Using `lsq-0/0/0`, CRTP can be applied on LFI and non-LFI packets. There will be no changes in their queue numbers because of CRTP.

Queuing on Q2s of Constituent Links

When using *class of service* on a multilink bundle, all Q2 traffic from the multilink bundle is queued to Q2 of constituent links based on a hash computed from the source address, destination address, and the IP protocol of the packet. If the IP payload is TCP or UDP traffic, the hash also includes the source port and destination port. As a result of this hash algorithm, all traffic belonging to one traffic flow is queued to Q2 of one constituent link. This method of traffic delivery to the constituent link is applied at all times, including when the bundle has not been set up with LFI.

Compressed Real-Time Transport Protocol Overview

Real-Time Transport Protocol (RTP) can help achieve interoperability among different implementations of network audio and video applications. However, in some cases, the header, which includes the IP, UDP, and RTP headers, can be too large (around 40 bytes) on networks using low-speed lines such as dial-up modems. Compressed Real-Time Transport Protocol (CRTP) can be configured to reduce network overhead on low-speed links. CRTP replaces the IP, UDP, and RTP headers with a 2-byte context ID (CID), reducing the header overhead considerably.

Figure 21 on page 345 shows how CRTP compresses the RTP header in a voice packet by reducing a 40-byte header to a 2-byte header.

Figure 21: CRTP

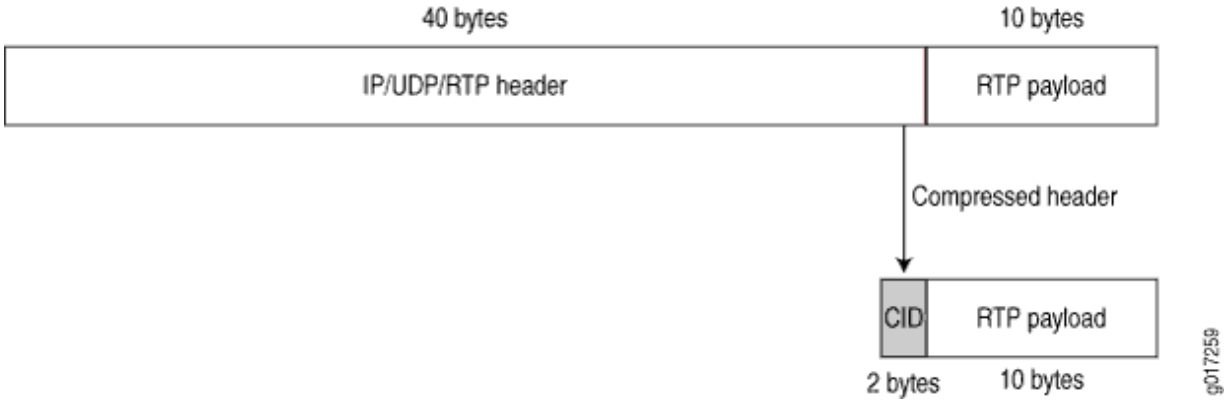

You can configure CRTP with MLPPP or PPP logical interface encapsulation on link services interfaces. See "Example: Configuring an MLPPP Bundle" on page 392.

Real-time and non-real-time data frames are carried together on lower-speed links without causing excessive delays to the real-time traffic. See "Understanding Link Fragmentation and Interleaving Configuration" on page 373.

Configuring Fragmentation by Forwarding Class

For `lsq-0/0/0`, you can specify fragmentation properties for specific forwarding classes. Traffic on each forwarding class can be either multilink encapsulated (fragmented and sequenced) or non-encapsulated (hashed with no fragmentation). By default, traffic in all forwarding classes is multilink encapsulated.

When you do not configure fragmentation properties for the queues on MLPPP interfaces, the fragmentation threshold you set at the `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`* `fragment-threshold]` hierarchy level is the fragmentation threshold for all forwarding classes within the MLPPP interface. For MLFR FRF.16 interfaces, the fragmentation threshold you set at the `[edit interfaces` *`interface-name`* `mlfr-uni-nni-bundle-options fragment-threshold]` hierarchy level is the fragmentation threshold for all forwarding classes within the MLFR FRF.16 interface.

If you do not set a maximum fragment size anywhere in the configuration, packets are still fragmented if they exceed the smallest maximum transmission unit (MTU) or maximum received reconstructed unit (MRRU) of all the links in the bundle. A non-encapsulated flow uses only one link. If the flow exceeds a single link, then the forwarding class must be multilink encapsulated, unless the packet size exceeds the MTU/MRRU.

Even if you do not set a maximum fragment size anywhere in the configuration, you can configure the MRRU by including the mrru statement at the `[edit interfaces` *`lsq-0/0/0`* `unit` *`logical-unit-number`*`]` or `[edit interfaces` *`interface-name`* `mlfr-uni-nni-bundle-options]` hierarchy level. The MRRU is similar to the MTU, but is specific to link services interfaces. By default the MRRU size is 1504 bytes, and you can configure it to be from 1500 through 4500 bytes.

To configure fragmentation properties on a queue, include the fragmentation-maps statement at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]
```

```
fragmentation-maps {
    map-name {
        forwarding-class class-name {
            fragment-threshold bytes;
            multilink-class number;
            no-fragmentation;
        }
    }
}
```

To set a per-forwarding class fragmentation threshold, include the `fragment-threshold` statement in the fragmentation map. This statement sets the maximum size of each multilink fragment.

To set traffic on a queue to be non-encapsulated rather than multilink encapsulated, include the `no-fragmentation` statement in the fragmentation map. This statement specifies that an extra fragmentation header is not prepended to the packets received on this queue and that static link load balancing is used to ensure in-order packet delivery.

For a given forwarding class, you can include either the `fragment-threshold` or `no-fragmentation` statement; they are mutually exclusive.

You use the `multilink-class` statement to map a forwarding class into a multiclass MLPPP. For a given forwarding class, you can include either the `multilink-class` or `no-fragmentation` statement; they are mutually exclusive.

To associate a fragmentation map with a multilink PPP interface or MLFR FRF.16 DLCI, include the `fragmentation-map` statement at the `[edit class-of-service interfaces` *`interface-name`* `unit` **logical-unit-number**`]` hierarchy level:

```
[edit class-of-service interfaces]
```

```
lsq-0/0/0 {
    unit logical-unit-number { # Multilink PPP
    fragmentation-map map-name;
    }
}
```

```
lsq-0/0/0:channel { # MLFR FRF.16
    unit logical-unit-number
    fragmentation-map map-name;
    }
}
```

Configuring Link-Layer Overhead

Link-layer overhead can cause packet drops on constituent links because of bit stuffing on serial links. Bit stuffing is used to prevent data from being interpreted as control information.

By default, 4 percent of the total bundle bandwidth is set aside for link-layer overhead. In most network environments, the average link-layer overhead is 1.6 percent. Therefore, we recommend 4 percent as a safeguard.

For `lsq-0/0/0` on Juniper Networks device, you can configure the percentage of bundle bandwidth to be set aside for link-layer overhead. To do this, include the link-layer-overhead statement:

```
link-layer-overhead percent;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* mlfr-uni-nni-bundle-options]
- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-routers *logical-router-name* interfaces *interface-name* unit *logical-unit-number*]

You can configure the value to be from 0 percent through 50 percent.

Link Services Configuration Overview

Before you begin:

- Install device hardware.
- Establish basic connectivity. See the Getting Started Guide for your device.
- Have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions. See "Understanding Interfaces" on page 2

Plan how you are going to use the link services interface on your network. See "Link Services Interfaces Overview" on page 340.

To configure link services on an interface, perform the following tasks:

1. Configure link fragmentation and interleaving (LFI). See "Example: Configuring Link Fragmentation and Interleaving" on page 374.
2. Configure classifiers and forwarding classes. See "Example: Defining Classifiers and Forwarding Classes" on page 377.
3. Configure scheduler maps. See "Understanding How to Define and Apply Scheduler Maps" on page 382.
4. Configure interface shaping rates. See "Example: Configuring Interface Shaping Rates" on page 389
5. Configure an MLPPP bundle. See "Example: Configuring an MLPPP Bundle" on page 392.
6. To configure MLFR, see "Example: Configuring Multilink Frame Relay FRF.15" on page 398 or "Example: Configuring Multilink Frame Relay FRF.16" on page 403

7. To configure CRTP, see "Example: Configuring the Compressed Real-Time Transport Protocol" on page 410

Verifying the Link Services Interface

IN THIS SECTION

- Verifying Link Services Interface Statistics | 349
- Verifying Link Services CoS Configuration | 352

Confirm that the configuration is working properly.

Verifying Link Services Interface Statistics

IN THIS SECTION

- Purpose | 349
- Action | 349

Purpose

Verify the link services interface statistics.

Action

The sample output provided in this section is based on the configurations provided in "Example: Configuring an MLPPP Bundle" on page 392. To verify that the constituent links are added to the bundle correctly and the packets are fragmented and transmitted correctly, take the following actions:

1. On device R0 and device R1, the two devices used in this example, configure MLPPP and LFI as described in "Example: Configuring an MLPPP Bundle" on page 392.
2. From the CLI, enter the `ping` command to verify that a connection is established between R0 and R1.
3. Transmit 10 data packets, 200 bytes each, from R0 to R1.

4. On R0, from the CLI, enter the `show interfaces` *interface-name* `statistics` command.

```
user@R0> show interfaces lsq-0/0/0 statistics detail
Physical interface: lsq-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 29, Generation: 135
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2006-06-23 11:36:23 PDT (03:38:43 ago)
  Statistics last cleared: 2006-06-23 15:13:12 PDT (00:01:54 ago)
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                 1820                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                   10                    0 pps
  ...
  Egress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 DATA                          10                   10                    0
    1 expedited-fo                   0                    0                    0
    2 VOICE                          0                    0                    0
    3 NC                             0                    0                    0

  Logical interface lsq-0/0/0.0 (Index 67) (SNMP ifIndex 41) (Generation 133)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Bandwidth: 16mbps
    Bundle options:
      ...
      Drop timer period              0
      Sequence number format         long (24 bits)
      Fragmentation threshold        128
      Links needed to sustain bundle  1
      Interleave fragments           Enabled
    Bundle errors:
      Packet drops                   0 (0 bytes)
      Fragment drops                 0 (0 bytes)
     ...
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             0          0             0            0
        Output:            20          0          1920            0
```

```
      Packets:
        Input :              0          0             0            0
        Output:             10          0          1820            0
    Link:
      se-1/0/0.0
        Input :              0          0             0            0
        Output:             10          0          1320            0
      se-1/0/1.0
        Input :              0          0             0            0
        Output:             10          0           600            0
...
      Destination: 10.0.0.9/24, Local: 10.0.0.10, Broadcast: Unspecified, Generation:144
```

This output shows a summary of interface information. Verify the following information:

- `Physical interface`—The physical interface is `Enabled`. If the interface is shown as `Disabled`, do either of the following:
  - In the CLI configuration editor, delete the `disable` statement at the `[edit interfaces` *`interface-name`*`]` level of the configuration hierarchy.
  - In the J-Web configuration editor, clear the `Disable` check box on the `Interfaces>`*`interface-name`* page.
- `Physical link`—The physical link is `Up`. A link state of `Down` indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- `Last flapped`—The `Last Flapped` time is an expected value. The `Last Flapped` time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- `Traffic statistics`—Number and rate of bytes and packets received and transmitted on the interface. Verify that the number of inbound and outbound bytes and packets match the expected throughput for the physical interface. To clear the statistics and see only new changes, use the `clear interfaces statistics` *`interface-name`* command.
- `Queue counters`—Name and number of queues are as configured. This sample output shows that 10 data packets were transmitted and no packets were dropped.
- `Logical interface`—Name of the multilink bundle you configured—`lsq-0/0/0.0`.
- `Bundle options`—Fragmentation threshold is correctly configured, and fragment interleaving is enabled.
- `Bundle errors`—Any packets and fragments dropped by the bundle.

- `Statistics`—The fragments and packets are received and transmitted correctly by the device. All references to traffic direction (input or output) are defined with respect to the device. Input fragments received by the device are assembled into input packets. Output packets are segmented into output fragments for transmission out of the device.

  In this example, 10 data packets of 200 bytes were transmitted. Because the fragmentation threshold is set to 128 bytes, all data packets were fragmented into two fragments. The sample output shows that 10 packets and 20 fragments were transmitted correctly.

- `Link`—The constituent links are added to this bundle and are receiving and transmitting fragments and packets correctly. The combined number of fragments transmitted on the constituent links must be equal to the number of fragments transmitted from the bundle. This sample output shows that the bundle transmitted 20 fragments and the two constituent links `se-1/0/0.0` and `se-1/0/1.0.0` correctly transmitted `10+10=20` fragments.

- `Destination` and `Local`—IP address of the remote side of the multilink bundle and the local side of the multilink bundle. This sample output shows that the destination address is the address on R1 and the local address is the address on R0.

Verifying Link Services CoS Configuration

IN THIS SECTION

- Purpose | 352
- Action | 352

Purpose

Verify CoS configurations on the link services interface.

Action

From the CLI, enter the following commands:

- `show class-of-service interface` *`interface-name`*
- `show class-of-service classifier name` *`classifier-name`*
- `show class-of-service scheduler-map` *`scheduler-map-name`*

The sample output provided in this section is based on the configurations provided in"Example: Configuring an MLPPP Bundle" on page 392.

```
user@R0> show class-of-service interface lsq-0/0/0
Physical interface: lsq-0/0/0, Index: 136
Queues supported: 8, Queues in use: 4
  Scheduler map: [default], Index: 2
  Input scheduler map: [default], Index: 3
  Chassis scheduler map: [default-chassis], Index: 4
 Logical interface: lsq-0/0/0.0, Index: 69
    Object              Name                    Type              Index
    Scheduler-map       s_map                   Output            16206
    Classifier          ipprec-compatibility    ip                   12
```

```
user@R0> show class-of-service interface ge-0/0/1
Physical interface: ge-0/0/1, Index: 140
          Queues supported: 8, Queues in use: 4
          Scheduler map: [default], Index: 2
          Input scheduler map: [default], Index: 3

         Logical interface: ge-0/0/1.0, Index: 68
         Object              Name                  Type         Index
         Classifier         classfy_input           ip           4330
```

```
user@R0> show class-of-service classifier name classify_input
Classifier: classfy_input, Code point type: inet-precedence, Index: 4330

  Code point       Forwarding class           Loss priority
    000                 DATA                      low
    010                 VOICE                     low
```

```
user@R0> show class-of-service scheduler-map s_map
Scheduler map: s_map, Index: 16206

  Scheduler: DATA, Forwarding class: DATA, Index: 3810
  Transmit rate: 49 percent, Rate Limit: none, Buffer size: 49 percent, Priority:low
    Drop profiles:
      Loss priority        Protocol         Index         Name
```

```
    Low                    any             1         [default-drop-profile]
    Medium low             any             1         [default-drop-profile]
    Medium high            any             1         [default-drop-profile]
    High                   any             1         [default-drop-profile]

Scheduler: VOICE, Forwarding class: VOICE, Index: 43363
Transmit rate: 50 percent, Rate Limit: none, Buffer size: 5 percent, Priority:high
  Drop profiles:
    Loss priority        Protocol          Index        Name
    Low                     any              1        [default-drop-profile]
    Medium low              any              1        [default-drop-profile]
    Medium high             any              1        [default-drop-profile]
    High                    any              1        [default-drop-profile]

Scheduler: NC, Forwarding class: NC, Index: 2435
Transmit rate: 1 percent, Rate Limit: none, Buffer size: 1 percent, Priority:high
  Drop profiles:
    Loss priority        Protocol         Index      Name
    Low                     any              1       [default-drop-profile]
    Medium low              any              1       [default-drop-profile]
    Medium high             any              1       [default-drop-profile]
    High                    any              1       [default-drop-profile]
```

These output examples show a summary of configured CoS components. Verify the following information:

- `Logical Interface`—Name of the multilink bundle and the CoS components applied to the bundle. The sample output shows that the multilink bundle is `lsq-0/0/0.0`, and the CoS scheduler-map `s_map` is applied to it.
- `Classifier`—Code points, forwarding classes, and loss priorities assigned to the classifier. The sample output shows that a default classifier, `ipprec-compatibility`, was applied to the `lsq-0/0/0` interface and the classifier `classify_input` was applied to the `ge-0/0/1` interface.
- `Scheduler`—Transmit rate, buffer size, priority, and loss priority assigned to each scheduler. The sample output displays the data, voice, and network control schedulers with all the configured values.

Understanding the Internal Interface LSQ-0/0/0 Configuration

The link services interface is an internal interface only. It is not associated with a physical medium or PIM. Within an SRX Series device, packets are routed to this interface for link bundling or compression.

It may be required that you upgrade your configuration to use the internal interface lsq-0/0/0 as the link services queuing interface instead of ls-0/0/0, which has been deprecated. You can also roll back your modified configuration to use ls-0/0/0.

Example: Upgrading from ls-0/0/0 to lsq-0/0/0 for Multilink Services

IN THIS SECTION

- Requirements | 355
- Overview | 355
- Configuration | 356
- Verification | 359

This example shows how to upgrade from ls-0/0/0 to lsq-0/0/0 (or to reverse the change) for multilink services.

Requirements

This procedure is only necessary if you are still using ls-0/0/0 instead of lsq-/0/0/0 or if you need to revert to the old interface.

Overview

In this example, you rename the link services internal interface from ls-0/0/0 to lsq-0/0/0 or vice versa. You rename all occurrences of ls-0/0/0 in the configuration to lsq-0/0/0 and configure the fragmentation map by adding no fragmentation. You specify no fragmentation after the name of queue 2, if queue 2 is configured, or after assured forwarding. You then attach the fragmentation map configured in the preceding step to lsq-0/0/0 and specify the unit number as 6 of the multilink bundle for which interleave fragments is configured.

Then you roll back the configuration from lsq-0/0/0 to ls-0/0/0. You rename all occurrences in the configuration from lsq-0/0/0 to ls-0/0/0. You delete the fragmentation map if it is configured under the [class-of-service] hierarchy and delete the fragmentation map if it is assigned to lsq-0/0/0. You can delete multilink-max-classes if it is configured for lsq-0/0/0 under the [interfaces] hierarchy. You then delete link-layer-overhead if it is configured for lsq-0/0/0 under the [interfaces] hierarchy.

If no fragmentation is configured on any forwarding class and the fragmentation map is assigned to lsq-0/0/0, then you configure interleave fragments for the ls-0/0/0 interface. Finally, you configure the classifier for LFI packets to refer to queue 2. (The ls-0/0/0 interface treats queue 2 as the LFI queue.)

Configuration

IN THIS SECTION

- Procedure | 356

Procedure

CLI Quick Configuration

To quickly upgrade from ls-0/0/0 to lsq-0/0/0 (or reverse the change), copy the following commands and paste them into the CLI:

```
For interfaces ls-0/0/0 to lsq-0/0/0
[edit]
rename interfaces ls-0/0/0 to lsq-0/0/0
set class-of-service fragmentation-maps map6 forwarding-class assured-forwarding no-fragmentation
set class-of-service interfaces lsq-0/0/0 unit 6 fragmentation-map map6
```

```
For interfaces lsq-0/0/0 to ls-0/0/0
[edit]
rename interfaces lsq-0/0/0 to ls-0/0/0
delete class-of-service fragmentation-maps map6
delete class-of-service interfaces lsq-0/0/0 unit 6 fragmentation-map map6
delete  interfaces lsq-0/0/0 unit 6 link-layer-overhead
delete  interfaces lsq-0/0/0:0 mlfr-uni-nni-bundle-options link-layer-overhead
set interfaces ls-0/0/0 unit 6 interleave-fragments
```

Step-by-Step Procedure

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

To upgrade from ls-0/0/0 to lsq-0/0/0 or to reverse that change:

1. Rename all the occurrences of ls-0/0/0 in the configuration.

```
[edit]
user@host# rename interfaces ls-0/0/0 to lsq-0/0/0
```

2. Configure the fragmentation map.

```
[edit class-of-service fragmentation-maps]
user@host# set map6 forwarding-class assured-forwarding no-fragmentation
```

3. Specify the unit number of the multilink bundle.

```
[edit class-of-service ]
user@host# set interfaces lsq-0/0/0 unit 6 fragmentation-map map6
```

4. Roll back the configuration for all occurrences in the configuration.

```
[edit]
user@host# rename interfaces lsq-0/0/0 to ls-0/0/0
```

5. Delete fragmentation map under class of service.

```
[edit]
user@host# delete class-of-service fragmentation-maps map6
```

6. Delete fragmentation map if it is assigned to the lsq-0/0/0 interface.

```
[edit class-of-service interfaces]
user@host# delete lsq-0/0/0 unit 6 fragmentation-map map6
```

7. Delete multilink max classes if it is configured for lsq-0/0/0.

NOTE: Multilink-max-classes is not supported and is most likely not configured.

8. Delete link-layer-overhead if it is configured for lsq-0/0/0.

```
[edit interfaces]
user@host# delete lsq-0/0/0 unit 6 link-layer-overhead
```

9. Delete link-layer-overhead if it is configured for lsq-0/0/0:0.

```
[edit interfaces]
user@host# delete lsq-0/0/0:0 mlfr-uni-nni-bundle-options link-layer-overhead
```

10. Configure interleave fragments for the ls-0/0/0 interface.

```
[edit interfaces]
user@host# set ls-0/0/0 unit 6 interleave-fragments
```

Results

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

```
[edit]
    user@host# show class-of-service
    interfaces {
    lsq-0/0/0 {
        unit 6 {
        fragmentation-map map6;
        }
    }
    }
    fragmentation-maps {
        map6 {
        forwarding-class {
        assured-forwarding {
            no-fragmentation;
            }
        }
```

```
    }
}
```

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

Verification

IN THIS SECTION

- Verifying Link Services Internal Interface ls-0/0/0 to lsq-0/0/0 | **359**

Confirm that the configuration is working properly.

Verifying Link Services Internal Interface ls-0/0/0 to lsq-0/0/0

Purpose

Verify the link services internal interface ls-0/0/0 changed to lsq-0/0/0.

Action

From operational mode, enter the `show class-of-service` command.

Troubleshooting the Link Services Interface

IN THIS SECTION

- Determine Which CoS Components Are Applied to the Constituent Links | **360**
- Determine What Causes Jitter and Latency on the Multilink Bundle | **362**
- Determine If LFI and Load Balancing Are Working Correctly | **363**
- Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device | **372**

To solve configuration problems on a link services interface:

Determine Which CoS Components Are Applied to the Constituent Links

IN THIS SECTION

- Problem | 360
- Solution | 360

Problem

Description

You are configuring a multilink bundle, but you also have traffic without MLPPP encapsulation passing through constituent links of the multilink bundle. Do you apply all CoS components to the constituent links, or is applying them to the multilink bundle enough?

Solution

You can apply a scheduler map to the multilink bundle and its constituent links. Although you can apply several CoS components with the scheduler map, configure only the ones that are required. We recommend that you keep the configuration on the constituent links simple to avoid unnecessary delay in transmission.

Table 2 shows the CoS components to be applied on a multilink bundle and its constituent links.

Table 37: CoS Components Applied on Multilink Bundles and Constituent Links

Cos Component	Multilink Bundle	Constituent Links	Explanation
Classifier	Yes	No	CoS classification takes place on the incoming side of the interface, not on the transmitting side, so no classifiers are needed on constituent links.

Table 37: CoS Components Applied on Multilink Bundles and Constituent Links *(Continued)*

Cos Component	Multilink Bundle	Constituent Links	Explanation
Forwarding class	Yes	No	Forwarding class is associated with a queue, and the queue is applied to the interface by a scheduler map. The queue assignment is predetermined on the constituent links. All packets from Q2 of the multilink bundle are assigned to Q2 of the constituent link, and packets from all the other queues are queued to Q0 of the constituent link.
Scheduler map	Yes	Yes	Apply scheduler maps on the multilink bundle and the constituent link as follows: • Transmit rate—Make sure that the relative order of the transmit rate configured on Q0 and Q2 is the same on the constituent links as on the multilink bundle. • Scheduler priority—Make sure that the relative order of the scheduler priority configured on Q0 and Q2 is the same on the constituent links as on the multilink bundle. • Buffer size—Because all non-LFI packets from the multilink bundle transit on Q0 of the constituent links, make sure that the buffer size on Q0 of the constituent links is large enough. • RED drop profile—Configure a RED drop profile on the multilink bundle only. Configuring the RED drop profile on the constituent links applies a back pressure mechanism that changes the buffer size and introduces variation. Because this behavior might cause fragment drops on the constituent links, make sure to leave the RED drop profile at the default settings on the constituent links.

Table 37: CoS Components Applied on Multilink Bundles and Constituent Links *(Continued)*

Cos Component	Multilink Bundle	Constituent Links	Explanation
Shaping rate for a per-unit scheduler or an interface-level scheduler	No	Yes	Because per-unit scheduling is applied only at the end point, apply this shaping rate to the constituent links only. Any configuration applied earlier is overwritten by the constituent link configuration.
Transmit-rate exact or queue-level shaping	Yes	No	The interface-level shaping applied on the constituent links overrides any shaping on the queue. Thus apply transmit-rate exact shaping on the multilink bundle only.
Rewrite rules	Yes	No	Rewrite bits are copied from the packet into the fragments automatically during fragmentation. Thus what you configure on the multilink bundle is carried on the fragments to the constituent links.
Virtual channel group	Yes	No	Virtual channel groups are identified through firewall filter rules that are applied on packets only before the multilink bundle. Thus you do not need to apply the virtual channel group configuration to the constituent links.

SEE ALSO

- Class of Service User Guide (Security Devices)

Determine What Causes Jitter and Latency on the Multilink Bundle

IN THIS SECTION

- Problem | 363
- Solution | 363

Problem

Description

To test jitter and latency, you send three streams of IP packets. All packets have the same IP precedence settings. After configuring LFI and CRTP, the latency increased even over a noncongested link. How can you reduce jitter and latency?

Solution

To reduce jitter and latency, do the following:

1. Make sure that you have configured a shaping rate on each constituent link.
2. Make sure that you have not configured a shaping rate on the link services interface.
3. Make sure that the configured shaping rate value is equal to the physical interface bandwidth.
4. If shaping rates are configured correctly, and jitter still persists, contact the Juniper Networks Technical Assistance Center (JTAC).

Determine If LFI and Load Balancing Are Working Correctly

IN THIS SECTION

- Problem | 363
- Solution | 364

Problem

Description

In this case, you have a single network that supports multiple services. The network transmits data and delay-sensitive voice traffic. After configuring MLPPP and LFI, make sure that voice packets are transmitted across the network with very little delay and jitter. How can you find out if voice packets are being treated as LFI packets and load balancing is performed correctly?

Solution

When LFI is enabled, data (non-LFI) packets are encapsulated with an MLPPP header and fragmented to packets of a specified size. The delay-sensitive, voice (LFI) packets are PPP-encapsulated and interleaved between data packet fragments. Queuing and load balancing are performed differently for LFI and non-LFI packets.

To verify that LFI is performed correctly, determine that packets are fragmented and encapsulated as configured. After you know whether a packet is treated as an LFI packet or a non-LFI packet, you can confirm whether the load balancing is performed correctly.

`Solution Scenario`—Suppose two Juniper Networks devices, R0 and R1, are connected by a multilink bundle `lsq-0/0/0.0` that aggregates two serial links, `se-1/0/0` and `se-1/0/1`. On R0 and R1, MLPPP and LFI are enabled on the link services interface and the fragmentation threshold is set to 128 bytes.

In this example, we used a packet generator to generate voice and data streams. You can use the packet capture feature to capture and analyze the packets on the incoming interface.

The following two data streams were sent on the multilink bundle:

- 100 data packets of 200 bytes (larger than the fragmentation threshold)
- 500 data packets of 60 bytes (smaller than the fragmentation threshold)

The following two voice streams were sent on the multilink bundle:

- 100 voice packets of 200 bytes from source port 100
- 300 voice packets of 200 bytes from source port 200

To confirm that LFI and load balancing are performed correctly:

NOTE: Only the significant portions of command output are displayed and described in this example.

1. Verify packet fragmentation. From operational mode, enter the `show interfaces lsq-0/0/0` command to check that large packets are fragmented correctly.

```
user@R0#> show interfaces lsq-0/0/0
Physical interface: lsq-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 29
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
```

```
  Last flapped   : 2006-08-01 10:45:13 PDT (2w0d 06:06 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface lsq-0/0/0.0 (Index 69) (SNMP ifIndex 42)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
    Bandwidth: 16mbps
    Statistics         Frames      fps       Bytes        bps
    Bundle:
      Fragments:
        Input :             0        0           0          0
        Output:          1100        0      118800          0
      Packets:
        Input :             0        0           0          0
        Output:          1000        0      112000          0
   ...
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 9.9.9/24, Local: 9.9.9.10
```

Meaning—The output shows a summary of packets transiting the device on the multilink bundle. Verify the following information on the multilink bundle:

- The total number of transiting packets = 1000
- The total number of transiting fragments=1100
- The number of data packets that were fragmented =100

The total number of packets sent (600 + 400) on the multilink bundle match the number of transiting packets (1000), indicating that no packets were dropped.

The number of transiting fragments exceeds the number of transiting packets by 100, indicating that 100 large data packets were correctly fragmented.

Corrective Action—If the packets are not fragmented correctly, check your fragmentation threshold configuration. Packets smaller than the specified fragmentation threshold are not fragmented.

2. Verify packet encapsulation. To find out whether a packet is treated as an LFI or non-LFI packet, determine its encapsulation type. LFI packets are PPP encapsulated, and non-LFI packets are encapsulated with both PPP and MLPPP. PPP and MLPPP encapsulations have different overheads resulting in different-sized packets. You can compare packet sizes to determine the encapsulation type.

A small unfragmented data packet contains a PPP header and a single MLPPP header. In a large fragmented data packet, the first fragment contains a PPP header and an MLPPP header, but the consecutive fragments contain only an MLPPP header.

PPP and MLPPP encapsulations add the following number of bytes to a packet:

- PPP encapsulation adds 7 bytes:

  4 bytes of header+2 bytes of frame check sequence (FCS)+1 byte that is idle or contains a flag

- MLPPP encapsulation adds between 6 and 8 bytes:

  4 bytes of PPP header+2 to 4 bytes of multilink header

Figure 2 shows the overhead added to PPP and MLPPP headers.

Figure 22: PPP and MLPPP Headers

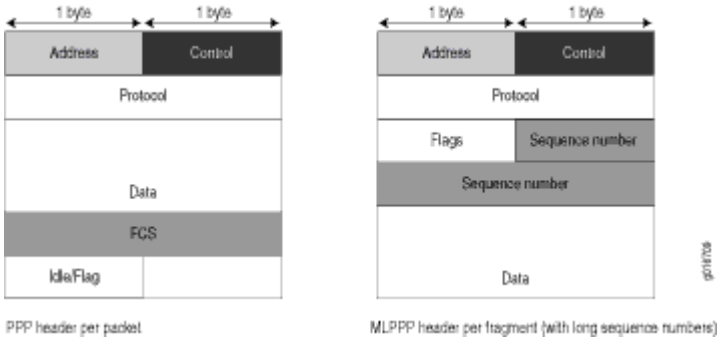

For CRTP packets, the encapsulation overhead and packet size are even smaller than for an LFI packet. For more information, see Example: Configuring the Compressed Real-Time Transport Protocol.

Table 3 shows the encapsulation overhead for a data packet and a voice packet of 70 bytes each. After encapsulation, the size of the data packet is larger than the size of the voice packet.

Table 38: PPP and MLPPP Encapsulation Overhead

Packet Type	Encapsulation	Initial Packet Size	Encapsulation Overhead	Packet Size after Encapsulation
Voice packet (LFI)	PPP	70 bytes	4 + 2 + 1 = 7 bytes	77 bytes

Table 38: PPP and MLPPP Encapsulation Overhead *(Continued)*

Packet Type	Encapsulation	Initial Packet Size	Encapsulation Overhead	Packet Size after Encapsulation
Data fragment (non-LFI) with short sequence	MLPPP	70 bytes	4 + 2 + 1 + 4 + 2 = 13 bytes	83 bytes
Data fragment (non-LFI) with long sequence	MLPPP	70 bytes	4 + 2 + 1 + 4 + 4 = 15 bytes	85 bytes

From operational mode, enter the `show interfaces queue` command to display the size of transmitted packet on each queue. Divide the number of bytes transmitted by the number of packets to obtain the size of the packets and determine the encapsulation type.

3. Verify load balancing. From operational mode, enter the `show interfaces queue` command on the multilink bundle and its constituent links to confirm whether load balancing is performed accordingly on the packets.

```
user@R0> show interfaces queue lsq-0/0/0
Physical interface: lsq-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 29
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: DATA
  Queued:
    Packets              :               600                 0 pps
    Bytes                :             44800                 0 bps
  Transmitted:
    Packets              :               600                 0 pps
    Bytes                :             44800                 0 bps
    Tail-dropped packets :                 0                 0 pps
    RED-dropped packets  :                 0                 0 pps
    …
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                 0                 0 pps
    Bytes                :                 0                 0 bps
```

```
  …
Queue: 2, Forwarding classes: VOICE
  Queued:
    Packets              :                400                 0 pps
    Bytes                :              61344                 0 bps
  Transmitted:
    Packets              :                400                 0 pps
    Bytes                :              61344                 0 bps
    …
Queue: 3, Forwarding classes: NC
  Queued:
    Packets              :                  0                 0 pps
    Bytes                :                  0                 0 bps
 …
```

```
user@R0> show interfaces queue se-1/0/0
Physical interface: se-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 35
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: DATA
  Queued:
    Packets              :                350                 0 pps
    Bytes                :              24350                 0 bps
  Transmitted:
    Packets              :                350                 0 pps
    Bytes                :              24350                 0 bps
   ...
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                  0                 0 pps
    Bytes                :                  0                 0 bps
  …
Queue: 2, Forwarding classes: VOICE
  Queued:
    Packets              :                100                 0 pps
    Bytes                :              15272                 0 bps
  Transmitted:
    Packets              :                100                 0 pps
    Bytes                :              15272                 0 bps
    …
```

```
Queue: 3, Forwarding classes: NC
  Queued:
    Packets              :                 19                0 pps
    Bytes                :                247                0 bps
  Transmitted:
    Packets              :                 19                0 pps
    Bytes                :                247                0 bps
    …
```

```
user@R0> show interfaces queue se-1/0/1
Physical interface: se-1/0/1, Enabled, Physical link is Up
  Interface index: 142, SNMP ifIndex: 38
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: DATA
  Queued:
    Packets              :                350                0 pps
    Bytes                :              24350                0 bps
  Transmitted:
    Packets              :                350                0 pps
    Bytes                :              24350                0 bps
    …
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              :                  0                0 pps
    Bytes                :                  0                0 bps
  …
Queue: 2, Forwarding classes: VOICE
  Queued:
    Packets              :                300                0 pps
    Bytes                :              45672                0 bps
  Transmitted:
    Packets              :                300                0 pps
    Bytes                :              45672                0 bps
    …
Queue: 3, Forwarding classes: NC
  Queued:
    Packets              :                 18                0 pps
    Bytes                :                234                0 bps
  Transmitted:
```

```
    Packets              :              18              0 pps
    Bytes                :             234              0 bps
```

Meaning—The output from these commands shows the packets transmitted and queued on each queue of the link services interface and its constituent links. Table 4 shows a summary of these values. (Because the number of transmitted packets equaled the number of queued packets on all the links, this table shows only the queued packets.)

Table 39: Number of Packets Transmitted on a Queue

Packets Queued	Bundle lsq-0/0/0.0	Constituent Link se-1/0/0	Constituent Link se-1/0/1	Explanation
Packets on Q0	600	350	350	The total number of packets transiting the constituent links (350+350 = 700) exceeded the number of packets queued (600) on the multilink bundle.
Packets on Q2	400	100	300	The total number of packets transiting the constituent links equaled the number of packets on the bundle.
Packets on Q3	0	19	18	The packets transiting Q3 of the constituent links are for keepalive messages exchanged between constituent links. Thus no packets were counted on Q3 of the bundle.

On the multilink bundle, verify the following:

- The number of packets queued matches the number transmitted. If the numbers match, no packets were dropped. If more packets were queued than were transmitted, packets were dropped because the buffer was too small. The buffer size on the constituent links controls congestion at the output stage. To correct this problem, increase the buffer size on the constituent links.
- The number of packets transiting Q0 (600) matches the number of large and small data packets received (100+500) on the multilink bundle. If the numbers match, all data packets correctly transited Q0.

- The number of packets transiting Q2 on the multilink bundle (400) matches the number of voice packets received on the multilink bundle. If the numbers match, all voice LFI packets correctly transited Q2.

On the constituent links, verify the following:

- The total number of packets transiting Q0 (350+350) matches the number of data packets and data fragments (500+200). If the numbers match, all the data packets after fragmentation correctly transited Q0 of the constituent links.

  Packets transited both constituent links, indicating that load balancing was correctly performed on non-LFI packets.

- The total number of packets transiting Q2 (300+100) on constituent links matches the number of voice packets received (400) on the multilink bundle. If the numbers match, all voice LFI packets correctly transited Q2.

  LFI packets from source port `100` transited `se-1/0/0`, and LFI packets from source port `200` transited `se-1/0/1`. Thus all LFI (Q2) packets were hashed based on the source port and correctly transited both constituent links.

`Corrective Action`—If the packets transited only one link, take the following steps to resolve the problem:

a. Determine whether the physical link is `up` (operational) or `down` (unavailable). An unavailable link indicates a problem with the PIM, interface port, or physical connection (link-layer errors). If the link is operational, move to the next step.

b. Verify that the classifiers are correctly defined for non-LFI packets. Make sure that non-LFI packets are not configured to be queued to Q2. All packets queued to Q2 are treated as LFI packets.

c. Verify that at least one of the following values is different in the LFI packets: source address, destination address, IP protocol, source port, or destination port. If the same values are configured for all LFI packets, the packets are all hashed to the same flow and transit the same link.

4. Use the results to verify load balancing.

Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device

IN THIS SECTION

- Problem | 372
- Solution | 372

Problem

Description

You are configuring a permanent virtual circuit (PVC) between T1, E1, T3, or E3 interfaces on a Juniper Networks device and a third-party device, and packets are being dropped and ping fails.

Solution

If the third-party device does not have the same FRF.12 support as the Juniper Networks device or supports FRF.12 in a different way, the Juniper Networks device interface on the PVC might discard a fragmented packet containing FRF.12 headers and count it as a "Policed Discard."

As a workaround, configure multilink bundles on both peers, and configure fragmentation thresholds on the multilink bundles.

Configuring Link Fragmentation and Interleaving

IN THIS SECTION

- Understanding Link Fragmentation and Interleaving Configuration | 373
- Example: Configuring Link Fragmentation and Interleaving | 374

The factor that determines the order in which output interface transmits traffic from an output queue is the priority scheduling on a multilink bundle. The large packets using this multilink bundle, cause delay

for the small and delay-sensitive packets to reach their turn for transmission. This delay renders some slow links like, T1 and E1, useless for delay-sensitive traffic. Link fragmentation and interleaving (LFI) solves this problem. The topics below topics the LFI in detail and its configuration.

Understanding Link Fragmentation and Interleaving Configuration

As it does on any other interface, priority scheduling on a multilink bundle determines the order in which an output interface transmits traffic from an output queue. The queues are serviced in a weighted round-robin fashion. But when a queue containing large packets starts using the multilink bundle, small and delay-sensitive packets must wait their turn for transmission. Because of this delay, some slow links, such as T1 and E1, can become useless for delay-sensitive traffic.

Link fragmentation and interleaving (LFI) solves this problem. It reduces delay and *jitter* on links by fragmenting large packets and interleaving delay-sensitive packets with the resulting smaller packets for simultaneous transmission across multiple links of a multilink bundle.

Figure 23 on page 374 illustrates how LFI works. In this figure, device R0 and device R1 have LFI enabled. When device R0 receives large and small packets, such as data and voice packets, it divides them into two categories. All voice packets and any other packets configured to be treated as voice packets are categorized as LFI packets and transmitted without fragmentation or an MLPPP header. If CRTP is configured on the bundle, LFI packets are transmitted through CRTP processing. The remaining non-LFI (data) packets can be fragmented or unfragmented based on the configured fragmentation threshold. The packets larger than the fragmentation threshold are fragmented. An MLPPP header (containing a multilink sequence number) is added to all non-LFI packets, fragmented and unfragmented.

The fragmentation is performed according to the fragmentation threshold that you configure. For example, if you configure a fragmentation threshold of 128 bytes, all packets larger than 128 bytes are fragmented. When device R1 receives the packets, it sends the unfragmented voice packets immediately but buffers the packet fragments until it receives the last fragment for a packet. In this example, when device R1 receives fragment 5, it reassembles the fragments and transmits the whole packet.

The unfragmented data packets are treated as a single fragment. Thus device R1 does not buffer the unfragmented data packets and transmits them as it receives them.

Figure 23: LFI on a Services Router

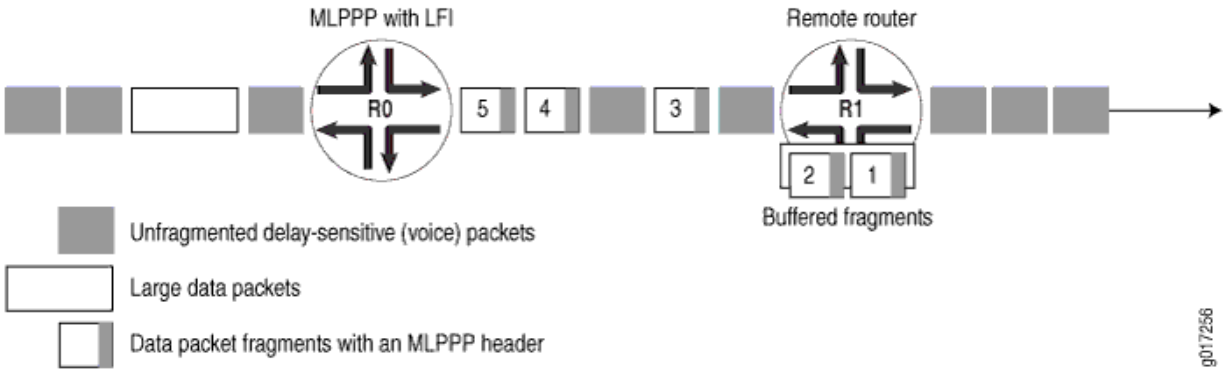

To configure LFI, you define the MLPPP encapsulation type and enable fragmentation and interleaving of packets by specifying the fragmentation threshold and fragmentation maps, with a no-fragmentation knob mapped to the forwarding class of choice.

Example: Configuring Link Fragmentation and Interleaving

IN THIS SECTION

- Requirements | 374
- Overview | 375
- Configuration | 375
- Verification | 376

This example shows how to configure LFI.

Requirements

Before you begin, you should have two Juniper Networks devices configured with at least two serial interfaces that communicate over serial links. This example shows two devices.

Overview

In this example, you create an interface called lsq-0/0/0. You specify the encapsulation type as multilink-ppp and set the fragmentation threshold value to 128. Set a fragmentation threshold of 128 bytes on the MLPPP bundle so that it applies to all traffic on both constituent links, enabling that any packet larger than 128 bytes transmitted on these links is fragmented. Any nonzero value must be a multiple of 64 bytes. The value can be between 128 and 16320. The default value is 0 bytes.

Configuration

IN THIS SECTION

- Procedure | 375

Procedure

Step-by-Step Procedure

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

To configure LFI:

1. Create an interface.

```
[edit]
user@host# edit interfaces lsq-0/0/0
```

2. Specify the encapsulation type and fragmentation threshold value.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 encapsulation multilink-ppp fragment-threshold 128
```

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

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

Verification

IN THIS SECTION

- Verifying Link Fragmentation and Interleaving Configuration | 376

Verifying Link Fragmentation and Interleaving Configuration

Purpose

Verify the LFI configuration.

Action

From operational mode, enter the `show interfaces lsq-0/0/0` command.

Configuring Class-of-Service on Link Services Interfaces

IN THIS SECTION

- Understanding How to Define Classifiers and Forwarding Classes | 377
- Example: Defining Classifiers and Forwarding Classes | 377
- Understanding How to Define and Apply Scheduler Maps | 382
- Example: Configuring Scheduler Maps | 384
- Understanding Interface Shaping Rates | 389
- Example: Configuring Interface Shaping Rates | 389

On a Juniper Networks device, when LFI is enabled, all forwarding traffic assigned to queue 2 or member link is treated as LFI (voice) traffic. The topics below discuss the overview of classifiers and

forwarding class, definition and application of schedule maps, and overview and configuration details of interface shaping rates on SRX series devices.

Understanding How to Define Classifiers and Forwarding Classes

By defining classifiers you associate incoming packets with a forwarding class and loss priority. Based on the associated forwarding class, you assign packets to output queues. To configure classifiers, you specify the bit pattern for the different types of traffic. The classifier takes this bit pattern and attempts to match it to the type of packet arriving on the interface. If the information in the packet's header matches the specified pattern, the packet is sent to the appropriate queue, defined by the forwarding class associated with the classifier.

On a Juniper Networks device, when LFI is enabled, all forwarding traffic assigned to queue 2 or member link is treated as LFI (voice) traffic. You do not need to assign network control traffic to a queue explicitly, because it is assigned to queue 3 by default.

NOTE: On member links:

- `DATA` is assigned to queue 0.
- `VOICE` is assigned to queue 2.
- `NC` (network control) is assigned to queue 3. By default NC is assigned to queue 3.

Example: Defining Classifiers and Forwarding Classes

IN THIS SECTION

- Requirements | **378**
- Overview | **378**
- Configuration | **378**
- Verification | **381**

This example shows how to define classifiers for different types of traffic, such as voice, data, and network control packets, and to direct the traffic to different output queues to manage your throughput.

Requirements

Before you begin:

- Configure two Juniper Networks devices with at least two serial interfaces that communicate over serial links.
- Configure CoS components. See *Junos OS Class of Service Configuration Guide for Security Devices*.

Overview

In this example, you configure class of service and set the default IP precedence classifier to classify_input, which is assigned to all incoming traffic. You then set the precedence bit value in the type of service field to 000 for all incoming data traffic and 010 for all incoming voice traffic. You set all outgoing data traffic to queue 0 and all voice traffic to queue 2, and fragmentation-map maps queue 2 to no fragmentation.

Configuration

IN THIS SECTION

- Procedure | 378

Procedure

CLI Quick Configuration

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

```
set class-of-service classifiers inet-precedence classify_input forwarding-class DATA loss-
priority low code-points 000
set class-of-service classifiers inet-precedence classify_input forwarding-class VOICE loss-
priority low code-points 010
set class-of-service forwarding-classes queue 0 DATA
set class-of-service forwarding-classes queue 2 VOICE
set class-of-service forwarding-classes queue 3 NC
```

```
set class-of-service interfaces ge-0/0/1 unit 0 classifiers inet-precedence classify_input
set class-of-service fragmentation-maps FM forwarding-class VOICE no-fragmentation
set class-of-service interfaces lsq-0/0/0 unit 0 fragmentation-map FM
```

Step-by-Step Procedure

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

To define classifiers and forwarding classes:

1. Configure class of service.

```
[edit]
user@host# edit class-of-service
```

2. Configure the behavior aggregate classifier for classifying packets.

```
[edit class-of-service]
user@host# edit classifiers inet-precedence classify_input
```

3. Assign packets with IP precedence to the data forwarding class and specify a loss priority.

```
[edit class-of-service classifiers inet-precedence classify_input]
user@host# set forwarding-class DATA loss-priority low code-points 000
```

4. Assign packets with IP precedence to the voice forwarding class and specify a loss priority.

```
[edit class-of-service classifiers inet-precedence classify_input]
user@host# set forwarding-class VOICE loss-priority low code-points 010
```

5. Specify the forwarding class one-to-one with the output queues.

```
[edit class-of-service]
user@host# edit forwarding-classes
user@host# set queue 0 DATA
```

```
user@host# set queue 2 VOICE
user@host# set queue 3 NC
```

6. Create an interface and apply the behavior aggregate classifier.

```
[edit class-of-service]
user@host# edit interfaces ge-0/0/1
user@host# set unit 0 classifiers inet-precedence classify_input
```

7. Configure fragmentation map.

```
[edit]
user@host# edit class-of-service
user@host# set fragmentation-maps FM forwarding-class VOICE no-fragmentation
```

8. Attach fragmentation map to the interface.

```
[edit class-of-service]
user@host# set interfaces lsq-0/0/0 unit 0 fragmentation-map FM
```

Results

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

```
[edit]
    user@host# show class-of-service
classifiers {
        inet-precedence classify_input {
            forwarding-class DATA {
        loss-priority low code-points 000;
        }
            forwarding-class VOICE {
        loss-priority low code-points 010;
        }
            }
        }
```

```
        forwarding-classes {
            queue 0 DATA;
            queue 2 VOICE;
            queue 3 NC;
            }
            interfaces {
                lsq-0/0/0 {
                unit 0 {
                fragmentation-map FM;
                    }
                    }
                    ge-0/0/1 {
                        unit 0 {
                        classifiers {
                    inet-precedence classify_input;
                    }
                }
            }
            }
            fragmentation-maps {
                FM {
                forwarding-class {
                VOICE {
                    no-fragmentation;
                }
            }
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying Classifiers and Forwarding Classes | **382**

To confirm that the configuration is working properly, perform this task:

Verifying Classifiers and Forwarding Classes

Purpose

Verify the classifiers and the forwarding classes.

Action

From operational mode, enter the `show class-of-service` command.

Understanding How to Define and Apply Scheduler Maps

Juniper Networks devices support per-unit scheduling `set class-of-service schedulers S0 priority low`, which allows you to configure scheduler maps on each MLPPP or MLFR multilink bundle. You can also configure scheduler maps on constituent links, but you must maintain the same relative priority on the constituent links and on the multilink bundle.

If you configure CoS components with LFI on a Juniper Networks device, we recommend that you follow certain recommendations for shaping rate, scheduling priority, and buffer size.

When you configure LFI, we recommend that you configure the shaping rate on each constituent link of the multilink bundle. Shaping rate configuration on the constituent links is required to limit the *jitter* on the LFI queue. If you anticipate no delay-sensitive or jitter-sensitive traffic on the LFI queue, or if there is no LFI traffic at all, shaping rate configuration is optional.

Table 40 on page 382 shows an example of correct and incorrect relative priorities on a multilink bundle and its constituent link. In this example, you have assigned a high priority to LFI packets and a low priority to data packets on the multilink bundle. To maintain the relative priority on the constituent links, you can assign a high priority to the LFI packets and a medium-high priority to the data packets, but you cannot assign a medium-high priority to LFI packets and a high priority to data packets.

Table 40: Relative Priorities on Multilink Bundles and Constituent Links

Multilink Bundle	Correct Constituent Link Priorities	Incorrect Constituent Link Priorities
LFI packets—High priority	LFI packets—High priority	LFI packet—Medium-high priority
Data packets—Low priority	Data packets—Medium-high priority	Data packets—High priority

By defining schedulers you configure the properties of output queues that determine the transmission service level for each queue. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, and the priority of the queue. After defining schedulers you associate them with forwarding classes by means of scheduler maps. You then associate each scheduler map with an interface, thereby configuring the hardware queues and packet schedulers that operate according to this mapping.

NOTE: When data and LFI streams are present, the following scheduler map configuration is recommended for constituent links. This gives less latency for LFI traffic and avoids out-of-order transmission of data traffic.

Configure the following schedulers:

- `set class-of-service schedulers S0 buffer-size temporal 20k`
- `set class-of-service schedulers S0 priority low`
- `set class-of-service schedulers S2 priority high`
- `set class-of-service schedulers S3 priority high`

Configure the following scheduler map:

- `set class-of-service scheduler-maps lsqlink_map forwarding-class best-effort scheduler S0`
- `set class-of-service scheduler-maps lsqlink_map forwarding-class assured-forwarding scheduler S2`
- `set class-of-service scheduler-maps lsqlink_map forwarding-class network-control scheduler S3`

Attach scheduler map to all member links:

- `set class-of-service interfaces t1-2/0/0 unit 0 scheduler-map lsqlink_map`

NOTE: Even after this configuration, if out-of-range sequence number drops are observed on the reassembly side, increase the drop-timeout of the bundle to 200 ms.

Example: Configuring Scheduler Maps

IN THIS SECTION

- Requirements | 384
- Overview | 384
- Configuration | 385
- Verification | 388

This example shows how to configure scheduler maps to determine the transmission service level for each output queue.

Requirements

Before you begin, you should have two Juniper Networks devices configured with at least two serial interfaces that communicate over serial links.

Overview

IN THIS SECTION

- Topology | 384

In this example, you create interfaces called lsq-0/0/0, se-1/0/0, and se-1/0/1. You enable per-unit scheduling to allow the configuration of scheduler maps on the bundle. You configure a scheduler map as s_map on lsq-0/0/0. You then apply the scheduler map to the constituent links, se-1/0/0 and se-1/0/1, of the multilink bundle. You associate the scheduler with each of the forwarding classes, DATA, VOICE and NC. You define the properties of output queues for the DATA scheduler by setting the transmit rate and the buffer size to 49 percent. You specify the properties of output queues for the VOICE scheduler by setting the transmit rate to 50 percent, the buffer size to 5 percent, and the priority to high. Finally, you define the properties of output queues for the NC scheduler by setting the transmit rate and the buffer size to 1 percent and the priority to high.

Topology

Configuration

IN THIS SECTION

- Procedure | 385

Procedure

CLI Quick Configuration

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

```
set interfaces lsq-0/0/0 per-unit-scheduler
set interfaces se-1/0/0 per-unit-scheduler
set interfaces se-1/0/1 per-unit-scheduler
set class-of-service interfaces lsq-0/0/0 unit 0 scheduler-map s_map
set class-of-service interfaces se-1/0/0 unit 0 scheduler-map s_map
set class-of-service interfaces se-1/0/1 unit 0 scheduler-map s_map
set class-of-service scheduler-maps s_map forwarding-class DATA scheduler DATA
set class-of-service scheduler-maps s_map forwarding-class VOICE scheduler VOICE
set class-of-service scheduler-maps s_map forwarding-class NC scheduler NC
set class-of-service schedulers DATA transmit-rate percent 49
set class-of-service schedulers DATA buffer-size percent 49
set class-of-service schedulers VOICE transmit-rate percent 50
set class-of-service schedulers VOICE buffer-size percent 5
set class-of-service schedulers VOICE priority high
set class-of-service schedulers NC transmit-rate percent 1
set class-of-service schedulers NC buffer-size percent 1
set class-of-service schedulers NC priority high
```

Step-by-Step Procedure

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

To configure scheduler maps:

1. Create interfaces and enable per-unit scheduling.

```
[edit interfaces]
user@host# set lsq-0/0/0 per-unit-scheduler
user@host# set se-1/0/0 per-unit-scheduler
user@host# set se-1/0/1 per-unit-scheduler
```

2. Define a scheduler map and apply it to the constituent links in the multilink bundle.

```
[edit class-of-service interfaces]
user@host# set lsq-0/0/0 unit 0 scheduler-map s_map
user@host# set se-1/0/0 unit 0 scheduler-map s_map
user@host# set se-1/0/1 unit 0 scheduler-map s_map
```

3. Associate a scheduler with each forwarding class.

```
[edit class-of-service scheduler-maps]
user@host# set s_map forwarding-class DATA scheduler DATA
user@host# set s_map forwarding-class VOICE scheduler VOICE
user@host# set s_map forwarding-class NC scheduler NC
```

4. Define the properties of output queues for the DATA scheduler.

```
[edit class-of-service schedulers]
user@host# set DATA transmit-rate percent 49
user@host# set DATA buffer-size percent 49
```

5. Define the properties of output queues for the VOICE scheduler.

```
[edit class-of-service schedulers]
user@host# set VOICE transmit-rate percent 50
user@host# set VOICE buffer-size percent 5
user@host# set VOICE priority high
```

6. Define the properties of output queues for the NC scheduler.

```
[edit class-of-service schedulers]
user@host# set NC transmit-rate percent 1
user@host# set NC buffer-size percent 1
user@host# set NC priority high
```

Results

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

```
[edit]
    user@host# show class-of-service
    interfaces {
    lsq-0/0/0 {
     unit 0 {
        scheduler-map s_map;
    }
    }
    se-1/0/0 {
        unit 0 {
        scheduler-map s_map;
    }
    }
    se-1/0/1 {
        unit 0 {
        scheduler-map s_map;
        }
    }
    }
        scheduler-maps {
    s_map {
    forwarding-class DATA scheduler DATA;
        forwarding-class VOICE scheduler VOICE;
        forwarding-class NC scheduler NC;
    }
}
    schedulers {
        DATA {
```

```
        transmit-rate percent 49;
        buffer-size percent 49;
        }
        VOICE {
            transmit-rate percent 50;
            buffer-size percent 5;
            priority high;
            }
            NC {
                transmit-rate percent 1;
                buffer-size percent 1;
            priority high;
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying the Configuration of scheduler maps. | 388

To confirm that the configuration is working properly, perform this task:

Verifying the Configuration of scheduler maps.

Purpose

Verify the configuration of scheduler maps.

Action

From operational mode, enter the `show class-of-services lsq-0/0/0 scheduler-map s_map`, `show class-of-services se-1/0/0 scheduler-map s_map`, and `show class-of-services se-1/0/1 scheduler-map s_map` commands.

Understanding Interface Shaping Rates

When you configure LFI, we recommend that you configure the shaping rate on each constituent link of the multilink bundle. Shaping rate configuration on the constituent links is required to limit the *jitter* on the LFI queue. If you anticipate no delay-sensitive or jitter-sensitive traffic on the LFI queue, or if there is no LFI traffic at all, shaping rate configuration is optional.

The shaping rate specifies the amount of bandwidth to be allocated for the multilink bundle. You must configure the shaping rate to be equal to the combined physical interface bandwidth for the constituent links. The combined bandwidth capacity of the two constituent links is 2 Mbps. Hence, configure a shaping rate of 2 Mbps on each constituent link.

Example: Configuring Interface Shaping Rates

IN THIS SECTION

- Requirements | **389**
- Overview | **390**
- Configuration | **390**
- Verification | **391**

This example shows how to configure interface shaping rates to control the maximum rate of traffic transmitted on an interface.

Requirements

Before you begin:

- Configure two Juniper Networks devices configured with at least two serial interfaces that communicate over serial links. For more information about serial interfaces. See Serial Interfaces Overview.
- To apply shaping rates to interfaces, you have to first enable per-unit scheduling. For more information on per-unit scheduling. See "Example: Configuring Scheduler Maps" on page 384.

Overview

IN THIS SECTION

- Topology | 390

In this example, you set the shaping rate to 2000000 for the constituent links of the multilink bundle, se-1/0/0 and se-1/0/1.

Topology

Configuration

IN THIS SECTION

- Procedure | 390

Procedure

Step-by-Step Procedure

To configure the interface shaping rates:

1. Configure class of service.

```
[edit]
user@host# edit class-of-service
```

2. Apply the shaping rates to the constituent links of the multilink bundle.

```
[edit class-of-service]
user@host# set interfaces se-1/0/0 unit 0 shaping-rate 2000000
user@host# set interfaces se-1/0/1 unit 0 shaping-rate 2000000
```

Verification

To verify the configuration is working properly, enter the `show class-of-service` command.

Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles

IN THIS SECTION

- Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links | **391**
- Example: Configuring an MLPPP Bundle | **392**

The topics below discuss the overview of MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links, and configuring an MLPP bundle on security devices.

Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links

Juniper Networks devices support MLPPP and MLFR multilink encapsulations. MLPPP multilink encapsulation enables you to bundle multiple PPP links into a single multilink bundle and MLFR multilink encapsulation enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1, E1, and serial links.

NOTE: Currently, Junos OS supports bundling of only one xDSL link under bundle interface.

You configure multilink bundles as logical units or channels on the link services interface `lsq-0/0/0`:

- With MLPPP and MLFR FRF.15, multilink bundles are configured as logical units on `lsq-0/0/0`—for example, `lsq-0/0/0.0` and `lsq-0/0/0.1`.

- With MLFR FRF.16, multilink bundles are configured as channels on `lsq-0/0/0`—for example, `lsq-0/0/0:0` and `lsq-0/0/0:1`.

After creating multilink bundles, you add constituent links to the bundle. The constituent links are the low-speed physical links that are to be aggregated. You can create 64 multilink bundles, and on each multilink bundle you can add up to 8 constituent links. The following rules apply when you add constituent links to a multilink bundle:

- On each multilink bundle, add only interfaces of the same type. For example, you can add either T1 or E1, but not both.
- Only interfaces with a PPP encapsulation can be added to an MLPPP bundle, and only interfaces with a Frame Relay encapsulation can be added to an MLFR bundle.
- If an interface is a member of an existing bundle and you add it to a new bundle, the interface is automatically deleted from the existing bundle and added to the new bundle.

Configuring a multilink bundle on the two serial links increases the bandwidth by 70 percent from approximately 1 Mbps to 1.7 Mbps and prepends each packet with a multilink header as specified in the FRF.12 standard. To increase the bandwidth further, you can add up to eight serial links to the bundle. In addition to a higher bandwidth, configuring the multilink bundle provides load balancing and redundancy. If one of the serial links fails, traffic continues to be transmitted on the other links without any interruption. In contrast, independent links require routing policies for load balancing and redundancy. Independent links also require IP addresses for each link as opposed to one IP address for the bundle. In the routing table, the multilink bundle is represented as a single interface.

Example: Configuring an MLPPP Bundle

IN THIS SECTION

- Requirements | **393**
- Overview | **393**
- Configuration | **393**
- Verification | **397**

This example shows how to configure an MLPPP bundle to increase traffic bandwidth.

Requirements

Before you begin, you should have two Juniper Networks devices configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you create the MLPPP bundle lsq-0/0/0.0 at the logical unit level of the link services interface lsq-0/0/0 on Juniper Networks devices R0 and R1. You then add the two serial interfaces se-1/0/0 and se-1/0/1 as constituent links to the multilink bundle. In Figure 24 on page 393, your company's branch office is connected to its main branch using devices R0 and R1. You transmit data and voice traffic on two low-speed 1-Mbps serial links. To increase bandwidth, you configure MLPPP and join the two serial links se-1/0/0 and se-1/0/1 into the multilink bundle lsq-0/0/0.0. Then you configure LFI and CoS on R0 and R1 to enable them to transmit voice packets ahead of data packets.

Figure 24: Configuring MLPPP and LFI on Serial Links

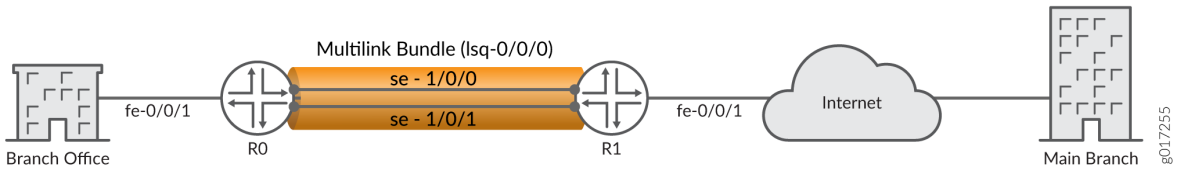

Configuration

IN THIS SECTION

- Procedure | **394**

Procedure

CLI Quick Configuration

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

```
For device R0
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.10/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
set interfaces se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz
```

```
For device R1
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.9/24
set interfaces se-1/0/0 unit 0 family mlppp bundle lsq-0/0/0.0
set interfaces se-1/0/1 unit 0 family mlppp bundle lsq-0/0/0.0
```

Step-by-Step Procedure

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

To configure MLPPP bundle:

1. Create an interface on both devices.

```
[edit]
user@host# edit interfaces lsq-0/0/0 unit 0
```

2. Configure a family inet and define the IP address on device R0.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.10/24
```

3. Configure a family inet and define the IP address on device R1.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.9/24
```

4. Specify the names of the constituent links to be added to the multilink bundle on both devices.

```
[edit interfaces]
user@host# edit se-1/0/0 unit 0
user@host# set family mlppp bundle lsq-0/0/0.0
[edit interfaces]
user@host# edit se-1/0/1 unit 0
user@host# set family mlppp bundle lsq-0/0/0.0
```

5. Set the serial options to the same values for both interfaces on R0.

> **NOTE**: R0 is set as a DCE device. The serial options are not set for interfaces on R1. You can set the serial options according to your network setup.

```
[edit interfaces]
user@host# set se-1/0/0 serial-options clocking-mode dce clock-rate 2.0mhz
user@host# set se-1/0/1 serial-options clocking-mode dce clock-rate 2.0mhz
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces lsq-0/0/0`, `show interfaces se-1/0/0`, and `show interfaces se-1/0/1` commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
    user@host# show interfaces lsq-0/0/0
family inet {
    address 10.0.0.10/24;
        }
    }
[edit]
```

```
user@host# show interfaces se-1/0/0
    clocking-mode dce;
        clock-rate 2.0mhz;
        }
        unit 0 {
            family mlppp {
            bundle lsq-0/0/0.0;
        }
    }
[edit]
    user@host# show interfaces se-1/0/1
serial-options {
    clocking-mode dce;
        clock-rate 2.0mhz;
        }
        unit 0 {
            family mlppp {
            bundle lsq-0/0/0.0;
        }
    }
```

```
For device R1
[edit]
user@host# show interfaces lsq-0/0/0
    family inet {
        address 10.0.0.9/24;
        }
    }
[edit]
    user@host# show interfaces se-1/0/0
    unit 0 {
    family mlppp {
        bundle lsq-0/0/0.0;
    }
}
[edit]
    user@host# show interfaces se-1/0/1
    unit 0 {
    family mlppp {
        bundle lsq-0/0/0.0;
```

```
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the MLPPP Bundle | **397**

Confirm that the configuration is working properly.

Verifying the MLPPP Bundle

Purpose

Verify that the constituent links are added to the bundle correctly.

Action

From operational mode, enter the `show interfaces lsq-0/0/0 statistics` command.

Configuring Multilink Frame Relay

IN THIS SECTION

- Understanding Multilink Frame Relay FRF.15 | **398**
- Example: Configuring Multilink Frame Relay FRF.15 | **398**
- Understanding Multilink Frame Relay FRF.16 | **402**
- Example: Configuring Multilink Frame Relay FRF.16 | **403**

The topics below discuss the overview and configuration details of Multilink Frame Relay FRF.15 and overview and configuration details of Multilink Frame Relay FRF.16 for security devices.

Understanding Multilink Frame Relay FRF.15

The link services intelligent queuing interface `lsq-0/0/0` supports Multilink Frame Relay end-to-end (MLFR FRF.15).

With MLFR FRF.15, multilink bundles are configured as logical units on the link services intelligent queuing interface, such as `lsq-0/0/0.0`. MLFR FRF.15 bundles combine multiple permanent virtual circuits (PVCs) into one aggregated virtual circuit (AVC). This process provides fragmentation over multiple PVCs on one end and reassembly of the AVC on the other end. You can configure LFI and CoS with MLFR in the same way that you configure them with MLPPP.

Example: Configuring Multilink Frame Relay FRF.15

IN THIS SECTION

- Requirements | **398**
- Overview | **398**
- Configuration | **399**
- Verification | **402**

This example shows how to configure MLFR FRF.15 for additional bandwidth, load balancing, and redundancy by aggregating low-speed links such as T1, E1, and serial links.

Requirements

Before you begin, you should have two Juniper Networks devices configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you aggregate two T1 links to create the MLFR FRF.15 bundle on two Juniper Networks devices, R0 and R1, and set the interface to lsq-0/0/0. You configure a logical unit on the lsq-0/0/0

interface and set the family type to inet with address 10.0.0.4/24. Then you configure an IP address for the multilink bundle on the unit level of the interface.

You define the multilink bundle as an MLFR FRF.15 bundle by specifying the MLFR end-to-end encapsulation type. You specify the names of the constituent links to be added to the multilink bundle as t1-2/0/0 and t1-2/0/1 and set the encapsulation type to frame relay. You then define R0 as a DCE device and R1 as a DTE device. You set the DLCI value to 100 (range is 16 through 1022). Finally, you set the multilink bundle to lsq-0/0/0.0.

Configuration

IN THIS SECTION

- Procedure | 399

Procedure

CLI Quick Configuration

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

```
For device R0
set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.4/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 dce
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

```
For device R1
 set interfaces lsq-0/0/0 unit 0 family inet address 10.0.0.5/24
set interfaces lsq-0/0/0 unit 0 encapsulation multilink-frame-relay-end-to-end
set interfaces t1-2/0/0 encapsulation frame-relay
set interfaces t1-2/0/1 encapsulation frame-relay
set interfaces lsq-0/0/0 unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

Step-by-Step Procedure

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

To configure the MLFR FRF.15 bundle:

1. Create an interface on both devices.

```
[edit]
user@host# edit interfaces lsq-0/0/0 unit 0
```

2. Set a logical unit on the interface and define the family type for devices R0 and R1.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set family inet address 10.0.0.4/24
user@host# set family inet address 10.0.0.5/24
```

3. Define the multilink bundle as an MLFR FRF.15 bundle.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set encapsulation multilink-frame-relay-end-to-end
```

4. Specify the names of the constituent links to be added to the multilink bundle.

```
[edit interfaces]
user@host# set t1-2/0/0 encapsulation frame-relay
user@host# set t1-2/0/1 encapsulation frame-relay
```

5. Define device R0 as a DCE device.

```
[edit interfaces]
user@host# edit lsq-0/0/0
user@host# set dce
```

6. Specify the DLCI as well as the multilink bundle to which the interface is to be added.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 100 family mlfr-end-to-end bundle lsq-0/0/0.0
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces lsq-0/0/0`, `show interfaces t1-2/0/0`, and `show interfaces t1-2/0/1` commands for R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
For device R0
[edit]
    user@host# show interfaces lsq-0/0/0
    dce;
unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
        dlci 100;
    family inet {
        address 10.0.0.4/24;
            }
        family mlfr-end-to-end {
        bundle lsq-0/0/0.0;
            }
        }
[edit]
    user@host#show interfaces t1-2/0/0
encapsulation frame-relay;
    [edit]
        user@host# show interfaces t1-2/0/1
    encapsulation frame-relay;
```

```
For device R1
[edit]
    user@host# show interfaces lsq-0/0/0
    unit 0 {
    encapsulation multilink-frame-relay-end-to-end;
        dlci 100;
        family inet {
```

```
    address 10.0.0.5/24;
    }
    family mlfr-end-to-end {
        bundle lsq-0/0/0.0;
    }
}
[edit]
    user@host# show interfaces t1-2/0/0
encapsulation frame-relay;
    [edit]
        user@host# show interfaces t1-2/0/1
    encapsulation frame-relay;
```

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

Verification

IN THIS SECTION

- Verifying the MLFR FRF.15 Configuration | 402

Confirm that the configuration is working properly.

Verifying the MLFR FRF.15 Configuration

Purpose

Verify the MLFR FRF.15 configuration.

Action

From operational mode, enter the `show interfaces` command.

Understanding Multilink Frame Relay FRF.16

The link services intelligent queuing interface `lsq-0/0/0` supports the Multilink Frame Relay (MLFR) user-to-network interface (UNI) and network-to-network interface (NNI) (MLFR FRF.16).

MLFR FRF.16 configures multilink bundles as channels on the link services intelligent queuing interface, such as `lsq-0/0/0:0`. A multilink bundle carries Frame Relay permanent virtual circuits (PVCs), identified by their data-link connection identifiers (DLCIs). Each DLCI is configured at the logical unit level of the link services intelligent queuing interface and is also referred as a *logical interface*. Packet fragmentation and reassembly occur on each virtual circuit. You can configure LFI and CoS with MLFR in the same way that you configure them with MLPPP.

Example: Configuring Multilink Frame Relay FRF.16

IN THIS SECTION

- Requirements | **403**
- Overview | **403**
- Configuration | **404**
- Verification | **409**

This example shows how to configure MLFR FRF.16 for additional bandwidth, load balancing, and redundancy.

Requirements

Before you begin, you should have two Juniper Networks devices configured with at least two serial interfaces that communicate over serial links.

Overview

IN THIS SECTION

- Topology | **404**

In this example, you aggregate two T1 interfaces to create an MLFR FRF.16 bundle on two Juniper Networks devices, R0 and R1. You configure the chassis interface and specify the number of MLFR FRF.16 bundles to be created on the interface. You then specify the channel to be configured as a

multilink bundle and create interface lsq-0/0/0:0. You set the multilink bundle as an MLFR FRF.16 bundle by specifying the MLFR UNI NNI encapsulation type.

Then you define R0 as a DCE device and R1 as a DTE device. You configure a logical unit on the multilink bundle lsq-0/0/0:0, and set the family type to inet. You then assign a DLCI of 400 and an IP address of 10.0.0.10/24 to the multilink bundle. You create the T1 interfaces, t1-2/0/0 and t1-2/0/1, that are to be added as constituent links to the multilink bundle and define the Frame Relay encapsulation type. Finally, you set the multilink bundle to lsq-0/0/0:0.

Topology

Configuration

IN THIS SECTION

- Procedure | 404

Procedure

CLI Quick Configuration

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

```
For device R0
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0:0 dce
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.10/24
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
For device R1
set chassis fpc 0 pic 0 mlfr-uni-nni-bundles 1
set interfaces lsq-0/0/0:0 encapsulation multilink-frame-relay-uni-nni
set interfaces lsq-0/0/0 unit 0 dlci 400 family inet address 10.0.0.9/24
```

```
set interfaces t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
set interfaces t1-2/0/0 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
set interfaces t1-2/0/1 unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

Step-by-Step Procedure

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

To configure an MLFR FRF.16 bundle:

1. Configure a chassis interface.

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

2. Specify the number of MLFR bundles.

```
[edit chassis]
user@host# set fpc 0 pic 0 mlfr-uni-nni-bundles 1
```

3. Create an interface.

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

4. Specify the MLFR encapsulation type.

```
[edit interfaces lsq-0/0/0:0]
user@host# set encapsulation multilink-frame-relay-uni-nni
```

5. Set device R0 as a DCE device.

```
[edit interfaces lsq-0/0/0:0]
user@host# set dce
```

6. Specify a logical unit on the multilink bundle and set the family type.

```
[edit interfaces lsq-0/0/0]
user@host# set unit 0 dlci 400 family inet address 10.0.0.10/24
```

7. Create the T1 interfaces and set the Frame Relay encapsulation.

```
[edit interfaces]
user@host# set t1-2/0/0 encapsulation multilink-frame-relay-uni-nni
user@host# set t1-2/0/1 encapsulation multilink-frame-relay-uni-nni
```

8. Specify the multilink bundle to which the interface is to be added as a constituent link on device R0.

```
[edit interfaces t1-2/0/0]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

9. Specify the multilink bundle to which the interface is to be added as a constituent link on device R1.

```
[edit interfaces t1-2/0/1]
user@host# set unit 0 family mlfr-uni-nni bundle lsq-0/0/0:0
```

Results

From configuration mode, confirm your configuration by entering the `show` commands for devices R0 and R1. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For device R0

```
[edit chassis]
user@host#show
fpc 0 {
    pic 0 {
        mlfr-uni-nni-bundles 1;
```

```
    }
}
```

```
[edit interfaces lsq-0/0/0:0]
user@host#show
dce;
encapsulation multilink-frame-relay-uni-nni;
```

```
[edit interfaces lsq-0/0/0]
user@host#show
unit 0 {
    dlci 400;
    family inet {
        address 10.0.0.10/24;
    }
}
```

```
[edit interfaces t1-2/0/0]
user@host#show
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle lsq-0/0/0:0;
    }
}
```

```
[edit interfaces t1-2/0/1]
user@host#show
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle lsq-0/0/0:0;
    }
}
```

For device R1

```
[edit chassis]
user@host#show
fpc 0 {
    pic 0 {
        mlfr-uni-nni-bundles 1;
    }
}
```

```
[edit interfaces lsq-0/0/0:0]
user@host#show
encapsulation multilink-frame-relay-uni-nni;
```

```
[edit interfaces t1-2/0/0]
user@host#show
encapsulation multilink-frame-relay-uni-nni;
unit 0 {
    family mlfr-uni-nni {
        bundle lsq-0/0/0:0;
    }
}
```

```
[edit interfaces t1-2/0/1]
user@host#show
encapsulation multilink-frame-relay-uni-nni;
    unit 0 {
        family mlfr-uni-nni {
    bundle lsq-0/0/0:0;
}
```

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

Verification

IN THIS SECTION

- Verifying the MLFR FRF.16 Configuration | 409

Confirm that the configuration is working properly.

Verifying the MLFR FRF.16 Configuration

Purpose

Verify the MLFR FRF.16 configuration.

Action

From operational mode, enter the `show interfaces` command.

Configuring Compressed Real-Time Transport Protocol

IN THIS SECTION

- Understanding Compressed Real-Time Transport Protocol | 410
- Example: Configuring the Compressed Real-Time Transport Protocol | 410

Compressed Real-Time Transport Protocol (CRTP) is typically used for compressing voice and video packets. The topics below discuss the overview of CRTP and its configuration details.

Understanding Compressed Real-Time Transport Protocol

Compressed Real-Time Transport Protocol (CRTP) is typically used for compressing voice and video packets. You can configure CRTP with LFI on a link services interface.

CRTP can be configured as a compression device on a T1 or E1 interface with PPP encapsulation, using the link services interface.

NOTE:

- F-max period—Maximum number of compressed packets allowed between transmission of full headers. It has a range from 1 to 65,535.
- Maximum and Minimum—UDP port values from 1 to 65,536 reserve these ports for RTP compression. CRTP is applied to network traffic on ports within this range. This feature is applicable only to voice services interfaces.

Example: Configuring the Compressed Real-Time Transport Protocol

IN THIS SECTION

- Requirements | **410**
- Overview | **411**
- Configuration | **411**
- Verification | **413**

This example shows how to configure CRTP to improve packet transmission, especially for time-sensitive voice packets.

Requirements

Before you begin, you should have two Juniper Networks devices configured with at least two serial interfaces that communicate over serial links.

Overview

In this example, you create a T1 interface called t1-1/0/0 and set the type of encapsulation to PPP. You set the link services intelligent queuing interface to lsq-0/0/0.0. You then create an interface called lsq-0/0/0 and set the logical unit 0. Finally, you set the F-max period to 2500, the minimum UDP port value to 2000, and the maximum UDP port value to 64009.

Configuration

IN THIS SECTION

- Procedure | 411

Procedure

CLI Quick Configuration

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

```
set interfaces t1-1/0/0 encapsulation ppp
set interfaces t1-1/0/0 unit 0 compression-device lsq-0/0/0.0
set interfaces lsq-0/0/0 unit 0 compression rtp f-max-period 2500 port minimum 2000 maximum
64009
```

Step-by-Step Procedure

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

To configure CRTP on a device:

1. Create the T1 interface.

```
[edit]
user@host# edit interfaces t1-1/0/0
```

2. Set the type of encapsulation.

```
[edit interfaces t1-1/0/0]
user@host# set encapsulation ppp
```

3. Add the link services intelligent queuing interface to the physical interface.

```
[edit interfaces t1-1/0/0]
user@host# edit unit 0
user@host# set compression-device lsq-0/0/0.0
```

4. Create an interface and set the logical unit.

```
[edit interfaces]
user@host# edit lsq-0/0/0 unit 0
```

5. Configure the link services intelligent queuing interface.

```
[edit interfaces lsq-0/0/0 unit 0]
user@host# set compression rtp f-max-period 2500 port minimum 2000 maximum 64009
```

Results

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

```
[edit]
    user@host# show interfaces
    lsq-0/0/0 {
    unit 0 {
        compression {
        rtp {
        f-max-period 2500;
            port minimum 2000 maximum 64009;
            }
        }
    }
```

```
}
    t1-1/0/0 {
        encapsulation ppp;
    unit 0 {
        compression-device lsq-0/0/0.0;
            }
        }
```

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

Verification

IN THIS SECTION

- Verifying the CRTP Configuration | **413**

Confirm that the configuration is working properly.

Verifying the CRTP Configuration

Purpose

Verify the CRTP configuration.

Action

From operational mode, enter the `show interfaces` command.

RELATED DOCUMENTATION

Link Services Interfaces Overview | **340**

CHAPTER 7

Configuring Management, Discard, and Loopback Interfaces

Configuring Management and Discard Interfaces | 415

Configuring Loopback Interfaces | 416

Configuring Management and Discard Interfaces

IN THIS SECTION

- Configuring Management Interfaces | 415
- Configuring Discard Interface | 416

The topics below discuss the over and configuration details of management and discard interfaces on the security devices.

Configuring Management Interfaces

Management interfaces are the primary interfaces for accessing the device remotely. Typically, a management interface is not connected to the in-band network, but is connected instead to the device's 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. SNMP can use the management interface to gather statistics from the device.

Management interfaces vary based on device type:

- The SRX5600 and SRX5800 devices include a 10/100-Mbps Ethernet port on the Routing Engine (RE). This port, which is labeled ETHERNET, is a dedicated out-of-band management interface for the device. Junos OS automatically creates the device's management interface `fxp0`. To use `fxp0` as a management port, you must configure its logical port `fxp0.0` with a valid IP address. While you can use `fxp0` to connect to a management network, you cannot place it into the management zone.

> **NOTE**: On the SRX5600 and SRX5800 devices, you must first connect to the device through the serial console port before assigning a unique IP address to the management interface.

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.

In an SRX Series device, the `fxp0` management interface is a dedicated port located on the Routing Engine. In an SRX Series chassis cluster configuration, the control link interface must be port `0` on an

SPC. For each node in the chassis cluster, you must configure the SPC that is used for the control link interface.

Configuring Discard Interface

The discard (dsc) interface is not a physical interface, but a virtual interface that discards packets. You can configure one discard interface. This interface allows you to identify the ingress (inbound) point of a denial-of-service (DoS) attack. When your network is under attack, the target host IP address is identified, and the local policy forwards attacking packets to the discard interface. Traffic routed out the discard interface is silently discarded.

RELATED DOCUMENTATION

Understanding Interfaces | 2

Configuring Loopback Interfaces

IN THIS SECTION

- Loopback Interface Overview | 416
- Configuring a Loopback Interface | 417

The topics below discuss the overview and configuration details of loopback interfaces on security devices.

Loopback Interface Overview

The Internet Protocol (IP) specifies a loopback network with the (IPv4) address 127.0.0.0/8. Most IP implementations support a loopback interface (lo0) to represent the loopback facility. Any traffic that a computer program sends on the loopback network is addressed to the same computer. The most

commonly used IP address on the loopback network is `127.0.0.1` for IPv4 and `::1` for IPv6. The standard domain name for the address is `localhost`.

A network device also includes an internal loopback interface (`lo0.16384`). The internal loopback interface is a particular instance of the loopback interface with the logical unit number 16384.

You use the loopback interface to identify the device. While you can use any interface address to determine if the device is online, the loopback address is the preferred method. Whereas interfaces might be removed or addresses changed based on network topology changes, the loopback address never changes.

When you ping an individual interface address, the results do not always indicate the health of the device. For example, a subnet mismatch in the configuration of two endpoints on a point-to-point link makes the link appear to be inoperable. Pinging the interface to determine whether the device is online provides a misleading result. An interface might be unavailable because of a problem unrelated to the device configuration or operation. You can use the loopback interface to address these issues.

Benefits

- As the loopback address never changes, it is the best way to identify a device in the network.
- The loopback interface is always up and reachable as long as the route to that IP address is available in the IP routing table. Hence, you can use the loopback interface for diagnostics and troubleshooting purposes.
- Protocols such as OSPF use the loopback address to determine protocol-specific properties for the device or network. Further, some commands such as `ping mpls` require a loopback address to function correctly.
- You can apply stateless firewall filters to the loopback logical unit to filter packets originating from, or destined for, the Routing Engine.
- Junos OS creates a separate loopback interface for the internal routing instance, which prevents any filter on `lo0.0` from disrupting internal traffic.

Configuring a Loopback Interface

The loopback interface supports many different network and operational functions and is an *always-up* interface. This means that the loopback interface ensures that the device is reachable, even if some of the physical interfaces are down or removed, or an IP address has changed. In most cases, you always define a loopback interface.

Junos OS follows the IP convention of identifying the loopback interface as lo0.

Junos OS requires that the loopback interface always be configured with a /32 network mask because the Routing Engine is essentially a host.

If you are using routing instances, you can configure the loopback interface for the default routing instance or for a specific routing instance. The following procedure adds the loopback interface to the default routing instance.

Optionally, instead of configuring the loopback interface at the `[edit interfaces]` hierarchy level, you can use a configuration group, as shown in this procedure. This is a recommended best practice for configuring the loopback interface. This procedure uses a group called `global` as an example.

To configure a loopback interface:

1. Using the host IP address, assign it to the loopback interface.

   Each host in your network deployment should have a unique loopback interface address. The address used here is only an example.

   ```
   [edit groups global interfaces lo0 unit 0 family inet]
   user@host# set address 192.0.2.27/32
   ```

2. (Optional) Set the preferred IP address.

   You can configure as many addresses as you need on the lo0 interface, so it is good practice to designate one preferred IP address.

   ```
   [edit groups global interfaces lo0 unit 0 family inet]
   user@host# set address 192.0.2.48/32 preferred
   ```

3. (Optional) Configure additional addresses.

   Only unit 0 is permitted as the primary loopback interface. If you want to add more IP addresses to unit 0, you configure them in the normal way under unit 0, without the `preferred` option.

   ```
   [edit groups global interfaces lo0 unit 0 family inet]
   user@host# set address 198.51.100.48/32
   user@host# set address 192.168.11.27
   ```

   NOTE: You do not have to include the `/32` as long as the IPv4 address is a valid host address. (This usually means that the last octet cannot be zero.)

4. Configure the localhost address.

On the lo0.0 interface, it is useful to have the IP address 127.0.0.1 configured, as certain processes such as NTP and MPLS ping use this default host address. The 127.0.0.1/32 address is a Martian IP address (an address invalid for routing), so it is never advertised by the Juniper Networks device.

```
[edit groups global interfaces lo0 unit 0 family inet]
user@host# set address 127.0.0.1/32
```

5. (Optional) Configure an ISO address.

   Depending on your network configuration, you might also need an ISO address for the IS-IS routing protocol.

```
[edit groups global interfaces lo0 unit 0 family iso]
user@host# address 49.0026.0000.0000.0110.00
```

6. If you used a configuration group, apply the configuration group, substituting `global` with the appropriate group name.

```
[edit]
user@host# set apply-groups global
```

7. Commit the configuration.

```
user@host# commit
```

8
CHAPTER

LTE Mini-PIM

LTE Mini Physical Interface Modules (LTE Mini-PIM) | 421

LTE Mini Physical Interface Modules (LTE Mini-PIM)

SUMMARY

Learn about the LTE MIni-PIM, the features supported on it and how to configure it on security devices.

IN THIS SECTION

- LTE Mini-PIM Overview | **421**
- Configure LTE Mini-PIM | **424**
- Example: Configure LTE Mini-PIM as a Backup Interface | **433**

LTE Mini-PIM Overview

IN THIS SECTION

- Features Supported on the LTE Mini-PIM | **422**

The LTE Mini-Physical Interface Module (Mini-PIM) provides wireless WAN support on security devices. Table 41 on page 421 specifies the key details of the LTE Mini-PIM interface.

Table 41: LTE Mini-PIM Device Details

Interface Details	Descriptions
Interface name	LTE Mini-PIM
Supported on	For information about platforms support, see hardware compatibility tool (HCT).
Models	• SRX-MP-LTE-AE • SRX-MP-LTE-AA See LTE Mini-PIM Models.

Table 41: LTE Mini-PIM Device Details *(Continued)*

Interface Details	Descriptions

Physical interface for the 4G LTE Mini-PIM	• The interface name is `cl-`*`slot number`*`/0/0` where *slot number* identifies the slot on the device in which you insert the LTE Mini-PIM. For example, `cl-1/0/0`.
	• Configurable properties on the physical interface are:
	  • A dialer pool to which the physical interface belongs and the priority of the interface in the pool.
	  • Profiles for the SIM cards.
	  • Radio access technology (automatic, 3G, LTE).

Key deployment	• Provides wireless WAN support. • Operates on 3G and 4G networks.

For hardware specifications for the LTE Mini-PIM, see LTE Mini-Physical Interface Module.

Features Supported on the LTE Mini-PIM

Table 42 on page 422 describes the key features supported on LTE Mini-PIM.

Table 42: Key Features Supported on LTE Mini-PIM

Feature	Description
Automatic switchover between service providers through dual SIMs	Supports dual Subscriber Identity Module (SIM) cards that allow connectivity to two different ISP networks. Automatic switchover provides a failover mechanism when the current active network fails.
Multiple service provider and Access Point Name (APN) profiles	Supports up to 16 profiles configuration for each SIM. The LTE Mini-PIM supports two SIM cards, you can configure a total of 32 profiles and at a time, only single profile is active.
LTE carrier aggregation	Supports increased capacity and network efficiency. Carrier aggregation expands the LTE bandwidth by combining secondary bands.

Table 42: Key Features Supported on LTE Mini-PIM *(Continued)*

Feature	Description
SIM security functions	Supports security functions such as SIM lock and unlock, and PIN change.
Primary, logical and backup interface with always-on, dial-on-demand, and backup modes	On receiving traffic, the logical dl0 interface enables and places calls through the physical interface in the dialer pool. The dialer interface performs backup and dialer filter functions. You can configure the dialer interface to operate as: • Primary Interface: The dialer interface connects to the network and is always on. For more information, see Configuring the LTE Mini-PIM as the Primary Interface. • Backup interface for the primary WAN connection: The dialer interface activates only when the primary connection fails. For more information, see Configuring the LTE Mini-PIM as a Backup Interface. • Dial-on-demand: For more information, see Configuring the LTE Interface as a Dial-on-Demand Interface. Configuration modes: always-on, dial-on-demand or backup modes. You can configure the Mini-PIM in any one of the modes. • Always-on: The Mini-PIM connects to the 3G/4G network after booting. The connection is always maintained. • When you configure as primary interface, the LTE Mini-PIM supports both the always-on and dial-on-demand modes.
Over-the-air upgrade for modem firmware	Supports Over-the-air (OTA) firmware upgrade that enables automatic and timely upgrade of modem firmware when new firmware versions are available. You can enable or disable the OTA upgrade on the LTE Mini-PIM. OTA upgrade is disabled by default.

Configure LTE Mini-PIM

IN THIS SECTION

- Configure LTE Mini-PIM as a Primary Interface | 424
- Configure LTE Mini-PIM in a High Availability Cluster Mode | 426
- Configure LTE Mini-PIM as a Backup Interface | 428
- Configure LTE Mini-PIM as a Dial-on-demand Interface | 430

You can configure the LTE Mini-PIM as a primary interface, as a backup interface or as a dial-on-demand interface.

Configure LTE Mini-PIM as a Primary Interface

Before you begin, ensure that dl0.0 is not configured as a backup interface. If dl0.0 is configured as a backup for any interface on the SRX Series device, then this configuration overrides the configuration outlined in this procedure, and the LTE Mini-PIM will function as a backup interface.

Use the `show interfaces | display set | match backup-option | match dl0.0` command to check whether any interface uses dl0.0 as a backup interface. If dl0.0 is configured as a backup interface, then delete the configuration by issuing the following command:

`delete interfaces` *`interface-name`* `unit 0 backup-options interface dl0.0`

The LTE Mini-PIM is installed on a SRX320 line of devices and functions as the primary interface as seen in Figure 1 and assumed that the LTE Mini-PIM is installed in slot 1 on the SRX320 line of devices.

Figure 25: LTE Mini-PIM as a Primary Interface

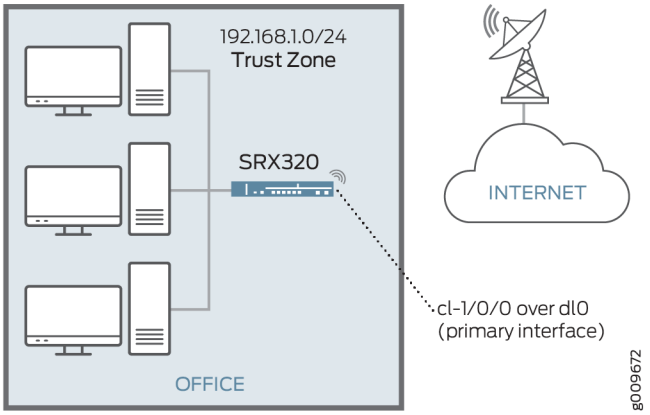

To configure the LTE Mini-PIM as a primary interface:

1. Configure the dialer interface.

```
user@host# set interfaces dl0 unit 0 family inet negotiate-address
user@host# set interfaces dl0 unit 0 family inet6 negotiate-address
user@host# set interfaces dl0 unit 0 dialer-options pool dialer-pool-number
user@host# set interfaces dl0 unit 0 dialer-options dial-string dial-number
user@host# set interfaces dl0 unit 0 dialer-options always-on
```

2. Configure the dialer pool for the LTE Mini-PIM physical interface.

```
user@host# set interfaces cl-1/0/0 dialer-options pool number
```

3. Configure the profile for the Subscriber Identity Module (SIM) cards.

```
user@host# run request modem wireless create-profile profile-id profile-id cl-1/0/0 slot sim-slot-number
access-point-name apn-name authentication-method none
```

`sim-slot-number` is the slot on the Mini-PIM in which the SIM card is inserted.

4. Verify that the profile is configured successfully.

```
user@host# run show modem wireless profiles cl-1/0/0 slot 1
```

5. Activate the SIM card.

```
user@host# set interfaces cl-1/0/0 act-sim sim-slot-number
```

6. Select the profile and configure the radio access type for the SIM card.

```
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number select-profile profile-id profile-id
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number radio-access automatic
```

   If a SIM card is installed in the second slot, then select the profile and configure the radio access type for the secondary SIM card.

7. Verify the status of the wireless network and dialer interface.

```
user@host#  run show modem wireless network
user@host#  run show interfaces dl0.0
```

   If the LTE Mini-PIM gets an IP address with a mask of /32 from the service provider, you can configure the default gateway information using the **set interfaces *cl-interface* cellular-options sim *sim-slot* gateway *ip-address/mask*** command to make the Mini-PIM accept the assigned IP address.

Configure LTE Mini-PIM in a High Availability Cluster Mode

An SRX chassis cluster supports two `cl` interfaces, `cl-1/1/0` (primary node) and `cl-8/1/0` (secondary node).

To configure the LTE Mini-PIM in a HA cluster mode:

1. Configure the dialer interface (`dl0`).

```
{primary:node0}[edit]
user@host# set interfaces dl0 unit 0 family inet negotiate-address
user@host# set interfaces dl0 unit 0 family inet6 negotiate-address
user@host# set interfaces dl0 unit 0 dialer-options pool dialer-pool-number
user@host# set interfaces dl0 unit 0 dialer-options dial-string dial-number
user@host# set interfaces dl0 unit 0 dialer-options always-on
```

2. Configure the LTE interface (`cl-1/1/0`) on the primary node.

a. Configure the dialer pool for the LTE physical interface.

```
{primary:node0}[edit]
user@host# set interfaces cl-1/1/0 dialer-options pool dialer-pool-number
```

b. Specify the priority for the interface. The interface with the higher priority becomes the active interface.

```
{primary:node0}[edit]
user@host# set interfaces cl-1/1/0 dialer-options pool dialer-pool-number priority priority
```

c. Configure the profile for the SIM cards.

```
{primary:node0}[edit]
user@host# run request modem wireless create-profile profile-id profile-id cl-1/1/0 slot
sim-slot-number access-point-name apn-name
```

d. Verify that the profile is configured successfully.

```
user@host# run show modem wireless profiles cl-1/1/0 slot 1
```

e. Activate the SIM card.

```
{primary:node0}[edit]
user@host# set interfaces cl-1/1/0 act-sim sim-slot-number
```

f. Select the profile and configure the radio access type for the SIM card.

```
{primary:node0}[edit]
user@host# set interfaces cl-1/1/0 cellular-options sim sim-slot-number select-profile
profile-id profile-id
user@host# set interfaces cl-1/1/0 cellular-options sim sim-slot-number radio-access
automatic
```

3. Repeat Step 2 to configure the LTE interface (`cl-8/1/0`) for the secondary node.

   If you assign the same priority to both interfaces, then the interface that is listed first in the configuration becomes the active interface.

Verify the active interface:

```
root@host> show dialer pools
Pool: 1
  Dialer interfaces:      Name            State
                          dl0.0           Active
  Subordinate interfaces: Name            Flags         Priority
                          cl-1/1/0        Active        100
                          cl-8/1/0        Inactive      1
```

4. Verify the status of the wireless network and dialer interface.

```
user@host# run show modem wireless network
user@host#  run show interfaces dl0.0
```

By default, the time interval taken to switch to the secondary cl interface when the active cl interface times out is 120 seconds. You can change the time interval by configuring the `redial-delay` option.

```
{primary:node0}[edit]
user@host# user@host# set interfaces dl0 unit 0 dialer-options redial-delay time-in-seconds
```

5. Verify that the profile is configured successfully.

```
user@host# run show modem wireless profiles cl-1/0/0 slot 1
```

6. Verify the status of the wireless network and dialer interface.

```
user@host#  run show modem wireless network
user@host#  run show interfaces dl0.0
```

Configure LTE Mini-PIM as a Backup Interface

You can configure the LTE Mini-PIM as a backup interface. If the primary interface fails, the LTE Mini-PIM connects to the network and remains online only until the primary interface becomes functional. The dialer interface is enabled only when the primary interface fails. LTE Mini-PIM installed on SRX320 and functions as a backup interface as shown in Figure 26 on page 429. The `ge-0/0/1` port is connected

to the internet and functions as the primary interface. In this scenario, the Mini-PIM is installed on slot 1.

Figure 26: LTE Mini-PIM as a Backup Interface

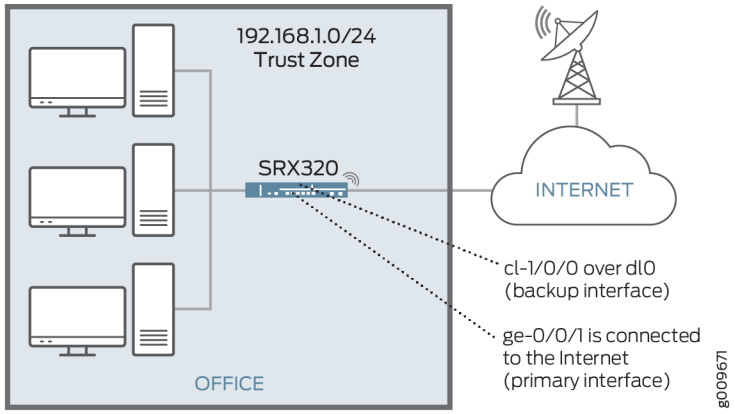

To configure the LTE Mini-PIM as a backup interface:

1. Configure the dialer interface.

```
user@host# set interfaces dl0 unit 0 family inet negotiate-address
user@host# set interfaces dl0 unit 0 family inet6 negotiate-address
user@host# set interfaces dl0 unit 0 dialer-options pool dialer-pool-number
user@host# set interfaces dl0 unit 0 dialer-options dial-string dial-number
```

2. Configure the dialer pool for the LTE Mini-PIM physical interface.

```
user@host# set interfaces cl-1/0/0 dialer-options pool dialer-pool-number
```

3. Configure the profile for the SIM cards.

   `sim-slot-number` is the slot on the Mini-PIM in which the SIM card is inserted.

```
user@host# run request modem wireless create-profile profile-id profile-id cl-1/0/0 slot sim-slot-number
access-point-name l3vpn.corp authentication-method none
```

4. Verify that the profile is configured successfully.

```
user@host# run show modem wireless profiles cl-1/0/0 slot 1
```

5. Activate the SIM card.

```
user@host# set interfaces cl-1/0/0 act-sim sim-slot-number
```

6. Select the profile and configure the radio access type for the SIM card.

```
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number select-profile profile-id profile-id
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number radio-access automatic
```

7. Configure the Ethernet interface as the primary interface, which connects to the wireless network. Configure the dl0 interface as the backup interface.

```
user@host# set interfaces ge-0/0/1 unit 0 family inet address 192.168.2.1/24
user@host# set interfaces ge-0/0/1 unit 0 backup-options interface dl0.0
```

8. Verify the status of the wireless network and dialer interface.

```
user@host#  run show modem wireless network
user@host#  run show interfaces dl0.0
```

You can use the `activation-delay` and `deactivation-delay` command-line options to avoid interface flaps. Avoid the Interface flaps by forcing a delay between the time the primary interface changes states, and the time the dialer interface is enabled or disabled. The activation delay controls the time between the primary interface going down and activation of the dialer interface. Similarly, the deactivation delay controls the time between the recovery of the primary interface and deactivation of the backup interface.

Configure LTE Mini-PIM as a Dial-on-demand Interface

When you configure the LTE interface as a primary interface, it functions either in always-on or in dial-on-demand mode. In always-on mode, the interface remains connected to the network whereas in dial-on-demand mode, the connection is established only when needed.

In dial-on-demand mode, you can enable the dialer interface only when network traffic configured as an “interesting traffic” arrives on the network. Interesting traffic triggers or activates the wireless WAN connection. You define an interesting packet by using the dialer filter. To configure dial-on-demand by using a dialer filter, you first configure the dialer filter and then apply the filter to the dialer interface. Once the traffic is sent over the network, an inactivity timer is triggered and the connection is closed after the timer expires. The dial-on-demand mode is supported only if the LTE Mini-PIM is configured as a primary interface.

The LTE Mini-PIM installed on an SRX320 functions as the primary interface as show in Figure 27 on page 431 and assumed that the LTE Mini-PIM is installed in slot 1 on the device.

Figure 27: LTE Mini-PIM as a Dial-on-Demand Interface

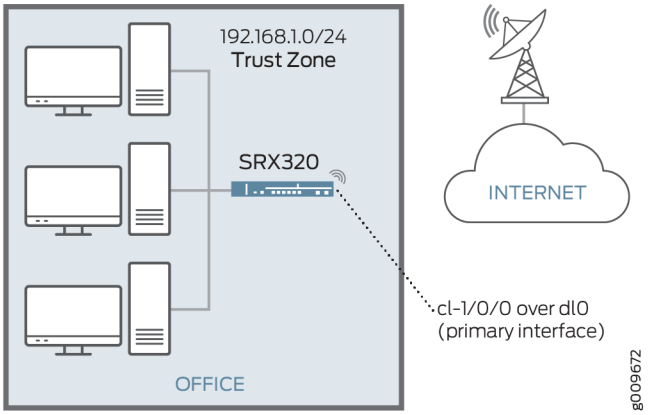

To configure the LTE Mini-PIM as a dial-on-demand interface:

1. Configure the dialer interface.

```
user@host# set interfaces dl0 unit 0 family inet negotiate-address
user@host# set interfaces dl0 unit 0 family inet6 negotiate-address
user@host# set interfaces dl0 unit 0 family inet filter dialer dialer-filter-name
user@host# set interfaces dl0 unit 0 dialer-options pool dialer-pool-number
user@host# set interfaces dl0 unit 0 dialer-options dial-string dial-number
```

   Optionally, you can configure the `idle-timeout` value, to determine the duration of the enabled connection in the absence of interesting traffic.

```
user@host# set interfaces dl0 unit 0 dialer-options idle-timeout idle-timeout-value
```

2. Configure the dialer pool for the LTE Mini-PIM physical interface.

```
user@host# set interfaces cl-1/0/0 dialer-options pool number
```

3. Create the dialer filter rule.

```
user@host# set firewall family inet dialer-filter dialer-filter-name term term1 from destination-
address ip-address then note
```

4. Set the default route.

```
set routing-options static route ip-address next-hop dl0.0
```

5. Configure the profile for the SIM cards.

```
user@host# run request modem wireless create-profile profile-id profile-id cl-1/0/0 slot sim-slot-
number access-point-name apn-name authentication-method none
```

6. Verify that the profile is configured successfully.

```
user@host# run show modem wireless profiles cl-1/0/0 slot 1
```

7. Activate the SIM card.

```
user@host# set interfaces cl-1/0/0 act-sim sim-slot-number
```

8. Select the profile and configure the radio access type for the SIM card.

```
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number select-profile profile-id
profile-id
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number radio-access automatic
```

9. Verify the configuration by sending traffic to the destination address. The traffic is routed to the dl0 interface and if it matches the dialer filter rule, then the dl0 is triggered to dial.
10. Verify the status of the wireless network and dialer interface.

```
user@host#  run show modem wireless network
user@host#  run show interfaces dl0.0
```

Example: Configure LTE Mini-PIM as a Backup Interface

IN THIS SECTION

- Requirements | 433
- Overview | 433
- Configuration | 433
- Verification | 436

This example shows how to configure the LTE Mini-PIM as a backup interface. If the primary interface fails, the Mini-PIM connects to the network and remains online only until the primary interface becomes functional. The dialer interface is enabled only when the primary interface fails. In this scenario, the Mini-PIM is installed on slot 1.

Requirements

Overview

Configuration

IN THIS SECTION

- CLI Quick Configuration | 433
- Configure the LTE Mini-PIM as a Backup Interface | 434
- Results | 435

CLI Quick Configuration

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

```
set interfaces dl0 unit 0 family inet negotiate-address
set interfaces dl0 unit 0 family inet6 negotiate-address
```

```
set interfaces dl0 unit 0 dialer-options pool dialer-pool-number
set interfaces dl0 unit 0 dialer-options dial-string dial-number
set interfaces cl-1/0/0 dialer-options pool dialer-pool-number
run request modem wireless create-profile profile-id profile-id cl-1/0/0 slot sim-slot-number
access-point-name l3vpn.corp authentication-method none
run show modem wireless profiles cl-1/0/0 slot 1
set interfaces cl-1/0/0 act-sim sim-slot-number
set interfaces cl-1/0/0 cellular-options sim sim-slot-number select-profile profile-id profile-id
set interfaces cl-1/0/0 cellular-options sim sim-slot-number radio-access automatic
set interfaces ge-0/0/1 unit 0 family inet address 192.168.2.1/24
set interfaces ge-0/0/1 unit 0 backup-options interface dl0.0
```

Configure the LTE Mini-PIM as a Backup Interface

Step-by-Step Procedure

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

To configure LTE Mini-PIM as a backup interface:

1. Create the dialer interface.

```
[edit interfaces]
user@host# set interfaces dl0 unit 0 family inet negotiate-address
user@host# set interfaces dl0 unit 0 family inet6 negotiate-address
user@host# set interfaces dl0 unit 0 dialer-options pool dialer-pool-number
user@host# set interfaces dl0 unit 0 dialer-options dial-string dial-number
```

2. Define the dialer pool for the LTE Mini-PIM physical interface.

```
user@host# set interfaces cl-1/0/0 dialer-options pool dialer-pool-number
```

3. Create and configure the profile on the SIM cards.

`sim-slot-number` is the slot on the Mini-PIM in which the SIM card is inserted.

```
user@host# run request modem wireless create-profile profile-id profile-id cl-1/0/0 slot sim-slot-number access-point-name l3vpn.corp authentication-method none
```

4. Activate the SIM card.

```
user@host# set interfaces cl-1/0/0 act-sim sim-slot-number
```

5. Select the profile and configure the radio access type for the SIM card.

```
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number select-profile profile-id profile-id
user@host# set interfaces cl-1/0/0 cellular-options sim sim-slot-number radio-access automatic
```

6. Specify Ethernet interface as the primary interface, which connects to the wireless network. Specify the dl0 interface as the backup interface.

```
user@host# set interfaces ge-0/0/1 unit 0 family inet address 192.168.2.1/24
user@host# set interfaces ge-0/0/1 unit 0 backup-options interface dl0.0
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces dl0.0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host>
Logical interface dl0.0 (Index 353) (SNMP ifIndex 559)
    Flags: Up Point-To-Point SNMP-Traps 0x4004000 Encapsulation: ENET2
    Dialer:
    State: Active, Dial pool: pool1
    Primary interface: ge-1/0/1.0 (Index 350)
    Dial strings: 1234
    Subordinate interfaces: cl-1/1/0 (Index 161)
    Activation delay: 0, Deactivation delay: 0
    Initial route check delay: 120
    Redial delay: 120
```

```
    Callback wait period: 5
    Load threshold: 0, Load interval: 60
    Input packets : 7
    Output packets: 10
    Protocol inet, MTU: 1490
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Flags: Sendbcast-pkt-to-re, Negotiate-Address
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 100.100.60.208/29, Local: 100.100.60.212, Broadcast: 100.100.60.215
    Protocol inet6, MTU: 1490
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop
cnt: 0
    Flags: Is-Primary, Negotiate-Address
    Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::5a00:bb0f:fcaa:7d00
```

Verification

IN THIS SECTION

- Verification of the configured profile | 436
- Verification of status of the dialer interface | 439
- Verification of status of the modem network and modem firmware | 440

Verification of the configured profile

Purpose

Verify that the profile is configured successfully.

Action

From operational mode, run the `show modem wireless profiles cl-1/0/0 slot 1` command.

```
user@host> show modem wireless profiles cl-1/0/0 slot 1

Profile details
  Max profiles: 16
```

```
  Default profile Id: 1

Profile 1: ACTIVE
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4V6
Profile 2: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 3: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 4: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 5: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 6: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 7: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 8: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 9: Inactive
```

```
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 10: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 11: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 12: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 13: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 14: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 15: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
Profile 16: Inactive
  Valid: TRUE
  Access point name (APN): airtelgprs.com
  Authentication: None
  IP Version: IPV4
```

Meaning

The output confirms the profile is active.

Verification of status of the dialer interface

Purpose

Verify that the dialer interface is configured successfully.

Action

From operational mode, run the `show interfaces dl0.0` command.

```
user@host> show interfaces dl0.0

    Logical interface dl0.0 (Index 353) (SNMP ifIndex 559)
    Flags: Up Point-To-Point SNMP-Traps 0x4004000 Encapsulation: ENET2
    Dialer:
    State: Active, Dial pool: pool1
    Primary interface: ge-1/0/1.0 (Index 350)
    Dial strings: 1234
    Subordinate interfaces: cl-1/1/0 (Index 161)
    Activation delay: 0, Deactivation delay: 0
    Initial route check delay: 120
    Redial delay: 120
    Callback wait period: 5
    Load threshold: 0, Load interval: 60
    Input packets : 7
    Output packets: 10
    Protocol inet, MTU: 1490
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Flags: Sendbcast-pkt-to-re, Negotiate-Address
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 100.100.60.208/29, Local: 100.100.60.212, Broadcast: 100.100.60.215
    Protocol inet6, MTU: 1490
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop
cnt: 0
    Flags: Is-Primary, Negotiate-Address
    Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::5a00:bb0f:fcaa:7d00
```

Meaning

The output confirms the interface `dlo` is configured and active.

Verification of status of the modem network and modem firmware

Purpose

Verify that the wireless network is configured, check the firmware, and check if the sim is active.

Action

From operational mode, enter the `show modem wireless network cl-1/0/0` command to verify the network status and `show modem wireless firmware cl-1/0/0` command to verify the firmware and sim status. Alternatively you can use the `show configuration` command to verify the complete status.

```
user@host> show modem wireless network cl-1/0/0

LTE Connection details
Connected time: 147
IP: 172.16.52.4
Gateway: 172.16.52.5
DNS: 123.123.123.123
Input bps: 0
Output bps: 0
Bytes Received: 1308
Bytes Transferred: 1164
Packets Received: 10
Packets Transferred: 10
Wireless Modem Network Info
Current Modem Status: Connected
Current Service Status: Normal
Current Service Type: PS
Current Service Mode: LTE
Current Band: B3
Network: UNICOM
Mobile Country Code (MCC): 460
Mobile Network Code (MNC): 1
Location Area Code (LAC): 65534
Routing Area Code (RAC): 0
Cell Identification: 4865903
```

```
Access Point Name (APN): abcde
Public Land Mobile Network (PLMN): CHN-UNICOM
Physical Cell ID (PCI): 333
International Mobile Subscriber Identification (IMSI): ***************
International Mobile Equipment Identification (IMEI/MEID): ***************
Integrate Circuit Card Identity (ICCID): 89860114721100697502
Reference Signal Receiving Power (RSRP): -97
Reference Signal Receiving Quality (RSRQ): -16
Signal to Interference-plus-Noise Ratio (SiNR): 0
Signal Noise Ratio (SNR): 0
Energy per Chip to Interference (ECIO): 0
```

Meaning

The output here shows the wireless modem network is connected and IP address of the fimware connected.

RELATED DOCUMENTATION

dialer-options | 574

9

CHAPTER

Wi-Fi MPIM

Wi-Fi Mini Physical Interface Module (MPIM) | 443

Wi-Fi Mini Physical Interface Module (MPIM)

IN THIS SECTION

- Wi-Fi Mini-Physical Interface Module Overview | 443
- Configure Wi-Fi Mini-PIM | 446

The Wi-Fi Mini-Physical Interface Module (Mini-PIM) for SRX Series devices provides an integrated wireless access point (or wireless LAN) solution along with routing, switching, and security in a single device. The topics below describes the overview and configuration of Wi-Fi Mini-PIM on SRX series devices.

Wi-Fi Mini-Physical Interface Module Overview

IN THIS SECTION

- Wireless LAN Interface in Chassis Cluster Mode | 444
- Wireless LAN Interface in Layer 3 (L3) Mode | 445
- Wireless LAN Interface in Layer 2 (L2) Mode | 445
- Features Supported on the Wi-Fi Mini-PIM | 445

Wi-Fi Mini-Physical Interface Module (Wi-Fi Mini-PIM) for SRX320, SRX340, SRX345, SRX380, and SRX550M provides an integrated wireless access point —or wireless LAN— along with routing, switching, and security in a single device. Mini-PIM supports the 802.11ac Wave 2 wireless standards and is backward compatible with 802.11a/b/g/n. You can use the three new models of the Wi-Fi Mini-PIM based on the regional wireless standard requirements;

- SRX-MP-WLAN-US — The model based on USA's wireless standard.
- SRX-MP-WLAN-IL — The model based on Israel's wireless standard.
- SRX-MP-WLAN-WW — The model for other countries.

You cannot change the country code for the SRX-MP-WLAN-US and SRX-MP-WLAN-IL models as they are fixed. The Wi-Fi Mini-PIM can coexist with other Mini-PIMs supported on the SRX Series device.Table 43 on page 445 provides a summary of the features supported on Mini-PIM.

Typical deployments for Wi-Fi Mini-PIM solution include:

- Secure wireless LAN connectivity to endpoint devices of corporate users at remote branch offices. 802.11ac, WPA2, 802.1X, and SSID-to-VLAN mapping features provide secure Wireless LAN connectivity.
- Direct network connectivity to the enterprise Internet of Things (IoT) devices. The security features on the SRX Series devices secure the IoT devices.

See How to Install the Wi-Fi Mini-PIM for SRX Series Services Gateways for more information about how to install the Wi-Fi Mini-PIM.

Wireless LAN Interface in Chassis Cluster Mode

The Mini-PIM is also supported in chassis cluster mode to provide redundancy. Wireless users are connected to the active interface in redundancy group. To support chassis cluster mode for wireless LAN interface Mini-PIM, you need to configure chassis cluster setup with two wireless LAN interfaces `wl-`*x*`/0/0` and `wl-`*y*`/0/0`, where *x* indicates the slot number which wireless LAN interface Mini-PIM plug in on the node 0 and *Y* indicates the slot number which wireless LAN interface Mini-PIM plug in on the node 1.

In chassis cluster mode, there is one wireless LAN interface active, the other wireless LAN interface is inactive. Wi-Fi client is associated to active wireless LAN interface.

Below are the list of events which trigger wireless LAN interface failover when:

- wireless LAN interface is abnormal.
- primary wireless LAN interface is down.
- Redundant group which wireless LAN interface belongs to failover manually.
- primary WLAN interface node is failed.

After wireless LAN interface failover, the original inactive wireless LAN interface is changed to active and the Wi-Fi client sessions are reconnected to the new primary wireless LAN interface.

With chassis cluster mode, WLAND process runs on both nodes. The WLAND on primary node pushes the WLAN configuration to PFE on two nodes, and then PFE forwards the configuration to local wireless LAN interface card so that two wireless LAN interface cards have the same configuration.

To monitor wireless LAN interface status, WLAND finds the wireless LAN interface to be abnormal, it can trigger redundant group failover. In Layer 3 mode, by default, wireless LAN interface activity

monitor is configured for WLAN high availability using the commands `set chassis cluster redundancy-group 1 interface-monitor wl-2/0/0 weight 255` and `set chassis cluster redundancy-group 1 interface-monitor wl-7/0/0 weight 255`.

The new primary wireless LAN interface is active and the abnormal wireless LAN interface card is restarted and goes to inactive state. The Wi-Fi client is reconnected to the active wireless LAN interface automatically since the configuration (radio, channel, bandwidth, ssid, and so on) on active WAP is same as the original wireless LAN interface.

Wireless LAN Interface in Layer 3 (L3) Mode

The interfaces are configured as subordinate interface of RETH using the command `set interfaces wl-x/0/0 gigether-options redundant-parent reth-interface`. You can add the RETH interface to one redundant group and set the priority for each node in the redundant group. Only one wireless LAN interface is active in the redundant group and the other one is inactive.

Wireless LAN Interface in Layer 2 (L2) Mode

You can build SRX devices in chassis cluster mode with wireless LAN interface Mini-PIM. The peer wireless LAN interfaces are configured in the same VLAN and the wireless LAN interface on the primary node of redundant group zero is chosen as active interface by default. L2 mode (`family ethernet-switching`) of wireless LAN interface behave like any other L2 switching port (trunk port).

Features Supported on the Wi-Fi Mini-PIM

Table 43 on page 445 lists the key features supported on the Wi-Fi Mini-PIM.

Table 43: Wi-Fi Mini-PIM Features

Feature	Description
2x2 MU-MIMO	Enables transmission of data to multiple clients simultaneously.
Dual radios	Both radios of 2.4 GHz and 5 GHz bands are simultaneously supported. The maximum supported speed is upto 1.2 Gbps.

Table 43: Wi-Fi Mini-PIM Features *(Continued)*

Feature	Description
Virtual access points (VAPs) and VLAN features	• Allows you to segment the WLAN into multiple broadcast domains that are the wireless equivalents of Ethernet VLANs. A single access point is segregated into multiple individual VAPs, simulating multiple access points in a single system. • An access point supports multiple VLANs, which can be distributed across VAPs and radios. • You can configure up to eight VAPs per radio. You can map up to 16 extended service set identifiers (ESSIDs) to individual VLANs. • The VLANs from the Mini-PIM software map to VLANs on Junos OS.
Co-existence of interfaces	The Wi-Fi Mini-PIM coexists with 4G LTE, VDSL, T1, and serial interfaces.
Client authentication methods	Client authentication methods supported are Wi-Fi Protected Access (WPA) Enterprise (WPA2 standards) and Wi-Fi Protected Access (WPA) Personal (AES-CCMP cipher suits and WPA2 standards).

Configure Wi-Fi Mini-PIM

IN THIS SECTION

- Configure Network Setting for the Wi-Fi Mini-PIM | **447**
- Configure VLANS | **452**
- Configure Multiple VLANs and SSIDs | **454**

You can configure the radios and virtual access points on the Wi-Fi Mini-PIM. This topic contains sections that describe the basic Wi-Fi Mini-PIM configuration at the wireless interface level. For more information about how to install a Wi-Fi Mini-PIM see How to Install the Wi-Fi Mini-PIM for SRX Series Services Gateways.

The following sections describe how to configure the Wi-Fi Mini-PIM on your SRX Series device.

Configure Network Setting for the Wi-Fi Mini-PIM

Configure wl- interface

The interface name for the Mini-PIM is denoted as `wl-x/0/0`, where *x* is the slot on the SRX Series Services Gateway in which the Mini-PIM is installed. The wl- interface is created automatically when you insert the Mini-PIM into the slot on the SRX Series device.

To configure the wireless LAN interface:

1. Define an interface for the Wi-Fi Mini-PIM.

```
[edit interfaces]
user@host# set interface wl-x/0/0 unit unit-number family inet address address
```

2. Configure the DHCP address pool.

```
[edit]
user@host# set access address-assignment pool pool-name  family inet network ip-prefix
user@host# set access address-assignment pool  pool-name family inet range  range-name  low
lower-limit-address
user@host# set access address-assignment pool pool-name family inet range  range-name high
upper-limit-address
user@host# set access address-assignment pool pool-name  family inet dhcp-attributes router
router-address
user@host# set access address-assignment pool pool-name family inet dhcp-attributes name-
server (address | hostname)
user@host# set system services dhcp-local-server group group-name interface wl-x/0/0.0
```

3. Configure the wireless interface to be part of a zone, assign required security policies and commit the configuration.

Configure access point

To configure the access point associated with the wireless LAN interface wl-*x*/0/0:

1. Configure the name of the wireless access point.

```
[edit]
user@host# set wlan access-point name interface wl-x/0/0
```

2. Set the country code (applicable only for SRX-MP-WLAN-WW models of the Mini-PIM).

NOTE: If you do not set the country code for the SRX-MP-WLAN-WW models, the Mini-PIM considers the country code as US. The country code for the SRX-MP-WLAN-US and SRX-MP-WLAN-IL models are set and cannot be changed.

```
[edit]
user@host# set wlan access-point name access-point-options country country-code
```

3. Set the physical location (location of your hardware device, example: 1st-floor).

```
[edit]
user@host# set wlan access-point name location location
```

4. Commit the configuration.

Configure Radios

Every access point has two radios—radio 1 operates at 5-GHz bandwidth and radio 2 operates at 2.4-GHz bandwidth. A VAP is configured based on the radio. You can configure up to eight VAPs per radio and map up to 16 ESSIDs to individual VLANs. Wi-Fi Mini-PIM supports both the radios (2.4 and 5 GHz) to work simultaneously. You can also disable a radio. Table 2 lists the modes supported on each radio.

Table 44: Supported Modes on Wi-Fi Mini-PIM Radios

Radio	Supported Modes
Radio 1 (5.0 GHz)	• an—802.11a and 802.11n clients operating on 5 GHz frequency can connect to the access point • acn—802.11a, 802.11n and 802.11ac clients operating on 5 GHz frequency can connect to the access point
Radio 2 (2.4 GHz)	• gn—802.11g, 802.11b and 802.11n clients operating in 2.4 GHz frequency can connect to the access point. This is the default mode for this radio. • g—802.11g clients operating in 2.4 GHz frequency can connect to the access point supported from Junos OS Release 20.4R1.

To configure the radio:

1. Configure the radio mode. Radio 1 supports acn and an modes. Radio 2 supports g and gn mode. Note that radio 1 operates at 5-GHz and radio 2 operates at 2.4-GHz.

```
For radio 1:
[edit]
user@host# set wlan access-point name radio 1 radio-options mode (an | acn)
```

```
For radio 2:
[edit]
user@host# set wlan access-point name radio 2 radio-options mode gn
```

2. Configure the channel number. If you select auto, then the Mini-PIM chooses the channel automatically. By default, channel number is set to `auto`.

```
[edit]
user@host# set wlan access-point name radio (1 | 2) radio-options channel number (auto /
channel-number)
```

3. Configure the channel bandwidth. The default channel bandwidth is 20 MHz for the 2.4 GHz radio and 40 MHz for the 5 GHz radio. You can only set 80 MHz as the channel bandwidth for 5 GHz radio and not for 2.4GHz

```
[edit]
user@host# set wlan access-point name radio (1| 2) radio-options channel bandwidth (20 | 40 |
80)
```

4. Configure the transmit power. You can configure the transmit power on a per-radio basis.

NOTE: When you configure the transmit power, the Mini-PIM card will fix transmit power to the specified value set, in this case, the power by rate functionality does not work. So it is recommended not to set transmit power to a specified value. When you do not configure the transmit power (do not fix the transmit power to a specified value), the power by rate functionality works. If you configure the transmit power percentage to 100, then it chooses

the option "auto", the behavior is similar to no transmit power configured and power by rate functionality will work.

```
[edit]
user@host# set wlan access-point name radio (1| 2) radio-options transmit-power percent
```

5. Commit the configuration.

   In countries where Dynamic Frequency Selection (DFS) is required, the Wi-Fi card performs appropriate checks for radar. DFS is enabled by default. If you set the `channel number` to `auto`, the access point selects the channel from the list of DFS and non-DFS channels. You can disable DFS by using the `dfs-off` option **set wlan access-point *name* radio 1 radio-options dfs-off**.

   Only the 5 GHz radio (radio 1) supports DFS.

   For more information on DFS, see Channels and Frequencies Supported on the Wi-Fi Mini-PIM.

Configure Virtual access Points (VAP)

VAPs allow segmentation of the wireless LAN into multiple broadcast domains that are the wireless equivalents of Ethernet VLANs. To configure the virtual access point:

1. Configure the VAP settings.

```
[edit]
user@host# set wlan access-point name radio (1| 2) virtual-access-point id description
description
```

```
[edit]
user@host# set wlan access-point name radio (1| 2) virtual-access-point id ssid ssid
```

2. Configure either the WPA Enterprise or the WPA Personal authentication methods for the VAP.

   none—The data transferred between clients and the access point is not encrypted. Clients can associate with the access point without any authentication.

```
[edit]
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security none
```

wpa-enterprise—The device authenticates through an 802.1X-compliant RADIUS server.

```
[edit]
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
enterprise cipher-suites ccmp
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
enterprise radius-port port
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
enterprise radius-key secret-key
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
enterprise wpa-version v2
```

wpa-personal—The device uses preshared keys (PSKs) or a passphrase for authentication and encryption.

```
[edit]
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
personal cipher-suites ccmp
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
personal key-type (ascii|hex)
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
personal key key
user@host# set wlan access-point name radio (1| 2) virtual-access-point id security wpa-
personal wpa-version v2
```

3. Configure and specify the upload and download rate limits on the Wi-Fi Mini-PIM. The range for `upload-limit` and `download-limit` is from 256 Kbps to 1,048,576 Kbps.

```
[edit]
user@host# set wlan access-point name radio (1| 2) virtual-access-point id upload-limit
upload-limit-rate
user@host# set wlan access-point name radio (1| 2) virtual-access-point id download-limit
download-limit-rate
```

4. Specify the number maximum number of clients that can be connected to the VAP.

```
[edit]
user@host# set wlan access-point name radio (1| 2) virtual-access-point id maximum-stations
number
```

5. Commit the configuration.

After completing the configuration successfully completed, you can view the parameters by using the `show wlan access-points name detail` command.

Configure VLANS

Configure VLANs based on VAP

(Optional) A single access point is segregated into multiple individual virtual access points (VAPs) simulating multiple access points in a single system. The access point supports multiple VLANs. To configure the VLAN ID based on the VAP:

1. Configure the VLAN for the wireless LAN interface (wl- interface). Follow the below steps to configure VLAN ID based on the VAP :

```
[edit]
user@host# set vlans vlan-name vlan-id vlan-id
user@host# set vlans vlan-name vlan-id-list vid-list
user@host# set interfaces wl-x/0/0 unit unit-number family ethernet-switching vlan members
all
```

2. Set trunk mode on the wl- interface.

```
[edit]
user@host# set interfaces wl-x/0/0 unit unit-number family ethernet-switching interface-mode
trunk
```

3. Set the native VLAN of the wl- interface.

```
[edit]
user@host# set interfaces wl-x/0/0 native-vlan-id vlan-id
```

When you configure native vlan, the wl- interface will add a tag when it receives an untagged packet and takes no action when it receives a tagged native-vlan-id packet.

4. Configure the access point for the wl- interface.

```
[edit]
user@host# set wlan access-point name interface wl-x/0/0
```

5. Configure all VAP parameters including the radio mode, channel number, and VAP SSID, VAP VLAN ID on the Wi-Fi Mini-PIM.

```
[edit]
user@host# set wlan access-point name radio (1| 2) radio-options mode (an| gn | acn)
user@host# set wlan access-point name radio (1| 2) radio-options channel number  (auto | channel-number)
user@host# set wlan access-point name radio (1| 2) virtual-access-point id ssid ssid
user@host# set wlan access-point name radio (1| 2) virtual-access-point id vlan vlan-id
```

6. Commit the configuration.

Configure WPA enterprise authentication

(Optional) Wi-Fi protected access (WPA) enterprise is Wi-Fi alliance standard that uses RADIUS server authentication with AES-CCMP cipher suite. With this mode you can use high security encryption along with a centrally managed user authentication. Only the WPA2 standard is supported. To configure the WPA enterprise authentication:

1. Configure the address book and assign a security zone.

```
[edit]
user@host# set security address-book book-name address address-name ip-prefix
user@host# set security address-book book-name attach zone trust
user@host# set security address-book book-name attach zone dot1x
```

2. Configure security source rule-set from trust zone to the WPA authentication.

```
[edit]
user@host# set security nat source rule-set rule-set-name from zone trust
user@host# set security nat source rule-set   rule-set-name to zone dot1x
```

3. Configure the security source to match the source and destination address.

```
[edit]
user@host# set security nat source rule-set  rule-set-name rule rule-name match source-address ip-address
user@host# set security nat source rule-set rule-set-name rule rule-name match destination-address ip-address
```

4. Configure the UDP protocol and security source on the interface.

```
[edit]
user@host# set security nat source rule-set rule-set-name rule rule-name match protocol udp
user@host# set security nat source rule-set rule-set-name rule rule-name then source-nat
interface
```

5. Assign the security policies to the source and destination address.

```
[edit]
user@host# set security policies from-zone trust to-zone dot1x policy internet-access match
source-address ip-address
user@host# set security policies from-zone trust to-zone dot1x policy internet-access match
destination-address ip-address
user@host# set security policies from-zone trust to-zone dot1x policy internet-access match
application any
user@host# set security policies from-zone trust to-zone dot1x policy internet-access then
permit
```

6. Commit the configuration.

After completing the configuration successfully completed, you can view the parameters by using the `show wlan access-points` *name* `virtual-access-points` command.

Configure Multiple VLANs and SSIDs

You can configure 8 VAPs on each radio and each VAP is identified by the SSID. Up to 16 SSIDs can be configured on the Wi-Fi Mini-PIM. You can map a VLAN to each SSID or you can assign a single VLAN for multiple SSIDs The client connects to the VAP using the SSID and is associated to the VLAN that is mapped to the SSID.

You can configure multiple SSIDs to provide varied levels of access to different devices and users. Here is a sample configuration for three different types of users connecting to different VAPs. Each VAP is associated with a different VLAN.

Interface	VLAN ID	Address pool	VAP	SSID	Address pool
wl-2/0/0.0	100	junosDHCPPool			192.168.2.0/24

(Continued)

Interface	VLAN ID	Address pool	VAP	SSID	Address pool
wl-2/0/0.10	10	junosDHCPPool1	VAP1	VAP-10	192.168.10.0/24
wl-2/0/0.20	20	junosDHCPPool2	VAP2	VAP-20	192.168.20.0/24
wl-2/0/0.30	30	junosDHCPPool3	VAP3	VAP-30	192.168.30.0/24

1. Configure the interface to be part of the security zone.

```
user@host# set interfaces wl-2/0/0 unit 0 vlan-id 100
user@host# set interfaces wl-2/0/0 unit 0 family inet address 192.168.2.1/24
```

2. Configure a security zone.

```
user@host# set wlan access-point name interface wl-2/0/0
user@host# set wlan access-point name access-point-options country US
user@host# set wlan access-point name location California
```

3. Enable the DHCP server on the interface and configure the address pool for the Wi-Fi interface:

```
user@host# set wlan access-point name radio 1 radio-options mode acn
user@host# set wlan access-point name radio 1 radio-options channel number auto
user@host# set wlan access-point name radio 1 radio-options channel bandwidth 40
```

4. Configure flexible VLAN tagging on the Wi-Fi interface:

```
user@host# set wlan access-point name radio 2 radio-options mode gn
user@host# set wlan access-point name radio 2 radio-options channel number auto
user@host# set wlan access-point name radio 2 radio-options channel bandwidth 40
```

5. Configure the VLANs

```
user@host# set wlan access-point name radio 1 virtual-access-point 1 description VAP1
user@host# set wlan access-point name radio 1 virtual-access-point 1 ssid VAP-10
user@host# set wlan access-point name radio 1 virtual-access-point 1 vlan 10
```

```
user@host# set wlan access-point name radio 1 virtual-access-point 1 security wpa-personal
cipher-suites ccmp
user@host# set wlan access-point name radio 1 virtual-access-point 1 security wpa-personal
key-type ascii
user@host# set wlan access-point name radio 1 virtual-access-point 1 security wpa-personal
key ascii-string
user@host# set wlan access-point name radio 1 virtual-access-point 1 security wpa-personal
wpa-version v2
user@host# set wlan access-point name radio 1 virtual-access-point 1 upload-limit 1000
user@host# set wlan access-point name radio 1 virtual-access-point 1 download-limit 1000
user@host# set wlan access-point name radio 1 virtual-access-point 1 maximum-stations 70
```

6. Repeat steps 2 through 5 for the wl-2/0/0.10, wl-2/0/0.20, and wl-2/0/0.30 interfaces.
7. Configure the access point settings:

```
user@host# set wlan access-point name radio 1 virtual-access-point 2 description VAP2
user@host# set wlan access-point name radio 1 virtual-access-point 2 ssid VAP-20
user@host# set wlan access-point name radio 1 virtual-access-point 2 vlan 20
user@host# set wlan access-point name radio 1 virtual-access-point 2 security wpa-personal
cipher-suites ccmp
user@host# set wlan access-point name radio 1 virtual-access-point 2 security wpa-personal
key-type ascii
user@host# set wlan access-point name radio 1 virtual-access-point 2 security wpa-personal
key ascii-string
user@host# set wlan access-point name radio 1 virtual-access-point 2 security wpa-personal
wpa-version v2
user@host# set wlan access-point name radio 1 virtual-access-point 2 upload-limit 1000
user@host# set wlan access-point name radio 1 virtual-access-point 2 download-limit 1000
user@host# set wlan access-point name radio 1 virtual-access-point 2 maximum-stations 80
```

8. Configure the radio settings:

   For radio 1:

```
user@host# set wlan access-point name radio 2 virtual-access-point 3 description VAP3
user@host# set wlan access-point name radio 2 virtual-access-point 3 ssid VAP-30
user@host# set wlan access-point name radio 2 virtual-access-point 3 vlan 30
user@host# set wlan access-point name radio 2 virtual-access-point 3 security wpa-personal
cipher-suites ccmp
user@host# set wlan access-point name radio 2 virtual-access-point 3 security wpa-personal
key-type ascii
user@host# set wlan access-point name radio 2 virtual-access-point 3 security wpa-personal
```

```
key ascii-string
user@host# set wlan access-point name radio 2 virtual-access-point 3 security wpa-personal
wpa-version v2
user@host# set wlan access-point name radio 2 virtual-access-point 3 upload-limit 1000
user@host# set wlan access-point name radio 2 virtual-access-point 3 download-limit 1000
user@host# set wlan access-point name radio 2 virtual-access-point 3 maximum-stations 70
```

For radio 2:

```
user@host# commit
```

9. Configure the VAPs.

VAP1:

```
user@host>  show wlan access-points
```

VAP2:

```
Active access points information

Access-Point   Type   Interface    Radio-mode/Channel
i03-22-ap      Int    wl-1/0/0     gn/2, an/157
```

VAP3:

```
user@host>  show wlan access-point i03-22-ap detail
```

10. Commit the configuration.

```
Active access point detail information
Access Point        : i03-22-ap
Type                : Internal
Location            : First Floor, Building 8
Serial Number       : 850001809
Firmware Version    : 10.1.3.8
Alternate Version   : 10.1.3.7
Country             : US
Access Interface    : wl-1/0/0
```

```
Packet Capture      : Disabled
Ethernet Port:
MAC Address         : 00:14:13:12:10:11
IPv4 Address        : 192.168.1.5
Radio1:
Status              : On
MAC Address         : 00:1F:12:E0:84:20
Mode                : IEEE 802.11a/n
Channel             : 124 (5620 MHz)
Radio2:
Status              : On
MAC Address         : 00:1F:12:E0:84:30
Mode                : IEEE 802.11g/n
Channel             : 3 (2422 MHz)
```

Verification

Display information about the parameters configured on the Wi-Fi Mini-PIM.

- To display the details of all the access points configured on the Mini-PIM:

```
user@host# set security zones security-zone trust interfaces wl-2/0/0.0
```

```
user@host# set security zones security-zone trust host-inbound-traffic system-services dhcp
```

- To display the status of the specific access point.

```
user@host# set system services dhcp-local-server group jdhcp-group interface wl-2/0/0.0
user@host# set access address-assignment pool junosDHCPPool family inet network 192.168.2.0/24
user@host# set access address-assignment pool junosDHCPPool family inet range junosRange low
192.168.2.2
user@host# set access address-assignment pool junosDHCPPool family inet range junosRange high
192.168.2.254
user@host# set access address-assignment pool junosDHCPPool family inet dhcp-attributes router
192.168.2.1
```

```
user@host# set interfaces wl-2/0/0 flexible-vlan-tagging
user@host# set interfaces wl-2/0/0 native-vlan-id 100
```

RELATED DOCUMENTATION

wlan | 767

10
CHAPTER

Interfaces Support for SRX100, SRX110, SRX210, SRX240, SRX550, SRX650, and SRX1400 Devices

Configuring 1-Port Clear Channel DS3/E3 GPIM | 461

Configuring 3G Wireless Modems for WAN Connections | 473

Configuring CDMA EV-DO Modem Cards | 494

Configuring USB Modems for Dial Backup | 501

Configuring DOCSIS Mini-PIM Interfaces | 527

Configuring 1-Port Clear Channel DS3/E3 GPIM

IN THIS SECTION

- Understanding the 1-Port Clear Channel DS3/E3 GPIM | **461**
- Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode | **466**
- Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode | **468**
- Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode | **471**

The 1-Port Clear Channel DS3/E3 GPIM is a channel interface that can support full-duplex DS3 (T3) or E3 line rates. The below topics shows the overview of the interface, examplse on how to configure the 1-Port Clear Channel DS3/E3 GPIM for DS3 port mode, E3 port mode and M23 mapping mode respectively.

Understanding the 1-Port Clear Channel DS3/E3 GPIM

IN THIS SECTION

- Supported Features | **462**
- Interface Naming | **462**
- Physical Interface Settings | **463**
- Logical Interface Settings | **463**

The 1-Port Clear Channel DS3/E3 Gigabit-Backplane *Physical Interface Module* (GPIM) for the device functions as a clear channel interface that can support full-duplex DS3 (T3) or E3 line rates of 44.796 or 34.368 Mbps, respectively. The DS3/E3 interface is a popular high-bandwidth WAN interface for large enterprise branch locations that enables high-quality voice, video, and data applications with reduced latency. The GPIM device does not support channelization, but it supports a subrate DS3/E3 configuration.

This topic includes the following sections:

Supported Features

The clear channel implementation provides such features as subrate and scrambling options used by major DSU vendors. The following key features are available depending on the interface and mode selections:

- Framed and unframed DS3 (default) and E3 port modes
- Support for frame relay, point-to-point, and HDLC serial encapsulation protocols
- Support for popular vendor algorithms for subrate and payload scrambling
- Support for generation and detection of loopback control codes (line-loopback activate and deactivate) and FEAC codes
- External and internal clocking support
- Support for DS3 and E3 network alarms
- Support for chassis clusters
- Support for anti-counterfeit check
- Loopback (local, remote, and payload) and BERT/PRBS/QRSS diagnostics support
- MTU size of 4474 bytes (default) and 9192 bytes (maximum)

Interface Naming

The following format represents the 1-Port Clear Channel DS3/E3 GPIM interface names:

```
type-fpc/pic/port
```

where:

- *type*—Media type (T3 or E3)
- *fpc*—Number of the Flexible PIC Concentrator (FPC) card on which the physical interface is located
- *pic*—Number of the PIC on which the physical interface is located
- *port*—Specific port on the PIC

Examples: t3-1/0/0 and e3-2/0/0

Physical Interface Settings

The 1-Port Clear Channel DS3/E3 GPIM supports IP configurations. Using the CLI, you can configure the 1-Port Clear Channel DS3/E3 GPIM to operate in either DS3 or E3 mode. By default, at installation the physical interface, t3-x/y/z, is enabled on the GPIM port operating in DS3 mode with T3 framing.

You can reset the mode of the physical interface to E3 using the edit chassis command:

```
[edit]
user@host# set chassis fpc 1 pic 0 port 0 framing e3
```

Logical Interface Settings

The *logical interface* for the device is determined by setting the t3-options or e3-options of the edit interfaces command.

You can specify the MTU size for the GPIM interface. Junos OS supports an MTU value of 4474 bytes for the default value or up to 9192 bytes for maximum jumbo GPIM implementations.

Table 11 on page 71 identifies network interface specifications for DS3 or E3 modes.

Table 45: 1-Port Clear Channel DS3/E3 GPIM Interface Options

Description	DS3 Mode	E3 Mode
Network Interface Specifications		
Line encoding	B3ZS	HDB3
Framing	• C-bit parity (default) • M23	G.751 (default)

Table 45: 1-Port Clear Channel DS3/E3 GPIM Interface Options *(Continued)*

Description	DS3 Mode	E3 Mode
Subrate and scrambling	Vendor algorithms supported: • Adtran • Digital Link • Kentrox • Larscom • Verilink	Vendor algorithms supported: • Digital Link • Kentrox
Network alarms	Supported in accordance with the ANSI specification: • Loss of signal (LOS) • Out of frame (OOF) • Loss of frame (LOF) • Alarm identification Signal (AIS) • Remote defect identification (RDI)	Supported in accordance with the ITU-T specification: • Loss of signal (LOS) • Out of frame (OOF) • Alarm identification signal (AIS) • Remote defect identification (RDI) • Phase- locked loop (PLL)

Table 45: 1-Port Clear Channel DS3/E3 GPIM Interface Options *(Continued)*

Description	DS3 Mode	E3 Mode
Error counters	Incremented during a periodic 1-second polling routine: • Line code violations (LCV) • P-bit code violations (PCV) • C-bit code violations (CCV) • Line errored seconds (LES) • P-bit errored seconds (PES) • C-bit errored seconds (CES) • Severely errored framing seconds (SEFS) • P-bit severely errored seconds (PSES) • C-bit severely errored seconds (CSES) • Unavailable seconds (UAS)	Incremented during a periodic 1-second polling routine: • Frame alignment error (FAE) • Bipolar coding violations (BCV) • Excessive zeros (EXZ) • Line code violations (LCV) • Line errored seconds (LES) • Severely errored framing seconds (SEFS) • Unavailable seconds (UAS)
HDLC Features		
MTU	Default (4474 bytes) or maximum jumbo (up to 9192 bytes)	Default (4474 bytes) or maximum jumbo (up to 9192 bytes)
Shared flag	Supported	Supported
Idle flag/fill (0x7e or all ones)	Supported	Supported
Counters	Runts, giants	Runts, giants

SEE ALSO

Interface Naming Conventions | 9

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode

IN THIS SECTION

- Requirements | 466
- Overview | 466
- Configuration | 467

This example configures the GPIM in the DS3 (T3) operation mode.

Requirements

Before you begin:

- Install the device as specified in the *SRX Series Services Physical Interface Modules Hardware Guide*.

Overview

IN THIS SECTION

- Topology | 466

This example configures the basic T3 interface and modifies the framing to C-bit parity mode.

Topology

Configuration

IN THIS SECTION

- Procedure | 467

Procedure

Step-by-Step Procedure

To configure the GPIM:

1. Verify the installation, location, and status of the GPIM. In this example, the GPIM is installed in slot 8/PIC 0 and is currently online.

```
user@host> show chassis fpc pic-status
```

```
Slot 0 Online FPC
   PIC 0 Online 4x GE Base PIC
Slot 2 Offline FPC
Slot 5 Offline FPC
Slot 6 Online FPC
   PIC 0 Online 4x CT1E1 gPIM
Slot 7 Offline FPC
Slot 8 Online FPC
   PIC 0 Online 1x CLR CH T3/E3
```

2. Set the IP address for the logical interface.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet address interface 192.107.1.230/24
```

3. Set the MTU value to 9018.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet mtu 9018
```

4. Set the framing mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options cbit-parity
```

5. Enable the unframed DS3 mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options unframed
```

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

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

7. To verify the configuration for your device, enter the following operational command:

```
user@host> show interfaces t3-8/0/0 extensive
```

Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode

IN THIS SECTION

- Requirements | 469
- Overview | 469
- Configuration | 469

This example modifies the default configuration for an E3 environment.

Requirements

Before you begin:

- Install the device as specified in the *SRX Series Services Physical Interface Modules Hardware Guide*.

Overview

IN THIS SECTION

- Topology | 469

This example configures the basic E3 interface.

Topology

Configuration

IN THIS SECTION

- Procedure | 469

Procedure

Step-by-Step Procedure

To configure the GPIM in E3 framing:

1. Verify the installation, location, and status of the GPIM. In this example, the GPIM is installed in slot 8/PIC 0 and is currently online.

```
user@host> show chassis fpc pic-status
```

```
Slot 0 Online FPC
   PIC 0 Online 4x GE Base PIC
Slot 2 Offline FPC
Slot 5 Offline FPC
Slot 6 Online FPC
   PIC 0 Online 4x CT1E1 gPIM
Slot 7 Offline FPC
Slot 8 Online FPC
   PIC 0 Online 1x CLR CH T3/E3
```

2. Change to E3 port mode.

```
[edit]
user@host# set chassis fpc 8 pic 0 port 0 framing e3
```

3. Reset the MTU value to 3474.

```
[edit]
user@host# set interfaces e3-8/0/0 unit 0 family inet mtu 3474
```

4. Enable the unframed mode.

```
[edit]
user@host# set interfaces e3-8/0/0 e3-options unframed
```

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

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

6. To verify the configuration for your device, enter the following operational command:

```
user@host> show interfaces e3-8/0/0 extensive
```

Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode

IN THIS SECTION

- Requirements | 471
- Overview | 471
- Configuration | 472

The following example configures the GPIM in DS3 with M23 mapping mode. Note that M23 mapping does not provide C-bit parity.

Requirements

Before you begin:

- Install the device as specified in the *SRX Series Services Physical Interface Modules Hardware Guide*.

Overview

IN THIS SECTION

- Topology | 471

This example configures the basic T3 interface and modifies the framing to M23 mode without C-bit parity.

Topology

Configuration

IN THIS SECTION

- Procedure | 472

Procedure

Step-by-Step Procedure

To configure the GPIM:

1. Verify the installation, location, and status of the GPIM. In this example, the GPIM is installed in slot 8/PIC 0 and is currently online.

```
user@host> show chassis fpc pic-status
```

```
Slot 0 Online FPC
   PIC 0 Online 4x GE Base PIC
Slot 2 Offline FPC
Slot 5 Offline FPC
Slot 6 Online FPC
   PIC 0 Online 4x CT1E1 gPIM
Slot 7 Offline FPC
Slot 8 Online FPC
   PIC 0 Online 1x CLR CH T3/E3
```

2. Set the IP address for the logical interface.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet address interface 192.107.1.230/24
```

3. Set the MTU value to 9018.

```
[edit]
user@host# set interfaces t3-8/0/0 unit 0 family inet mtu 9018
```

4. Set the framing mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options m23
```

5. Disable C-bit parity for M23 mode.

```
[edit]
user@host# set interfaces t3-8/0/0 t3-options no-cbit-parity
```

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

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

7. To verify the configuration for your device, enter the following operational command:

```
user@host> show interfaces t3-8/0/0 extensive
```

Configuring 3G Wireless Modems for WAN Connections

IN THIS SECTION

- 3G Wireless Modem Overview | 474
- 3G Wireless Modem Configuration Overview | 475

- Understanding the Dialer Interface | **477**
- Example: Configuring the Dialer Interface | **479**
- Understanding the 3G Wireless Modem Physical Interface | **488**
- Example: Configuring the 3G Wireless Modem Interface | **488**
- Understanding the GSM Profile | **490**
- Example: Configuring the GSM Profile | **491**
- Unlocking the GSM 3G Wireless Modem | **493**

The topics below discuss the overview and configuration of 3G Wireless Modem, dialer interface, and 3G Wireless Modem physical interface.

3G Wireless Modem Overview

3G refers to the third generation of mobile phone standards and technology based on the International Telecommunication Union (ITU) International Mobile Telecommunications-2000 (IMT-2000) global standard. 3G networks are wide area cellular telephone networks that have evolved to include high-data rate services of up to 3 Mbps. This increased bandwidth makes 3G networks a viable option as primary or backup wide area network (WAN) links for a branch office.

Juniper Networks security devices support 3G wireless interfaces (USB-based 3G modems). When used in a branch office, these devices can provide dial-out services to PC users and forward IP traffic through a service provider's cellular network.

Figure 28 on page 475 illustrates a basic setup for 3G wireless connectivity for two branch offices. Branch Office A has a T1 leased line as the primary wide area network (WAN) link and a 3G wireless

modem connection as the failover link. Branch Office B uses the 3G wireless modem connection as the primary WAN link.

Figure 28: Wireless WAN Connections for Branch Offices

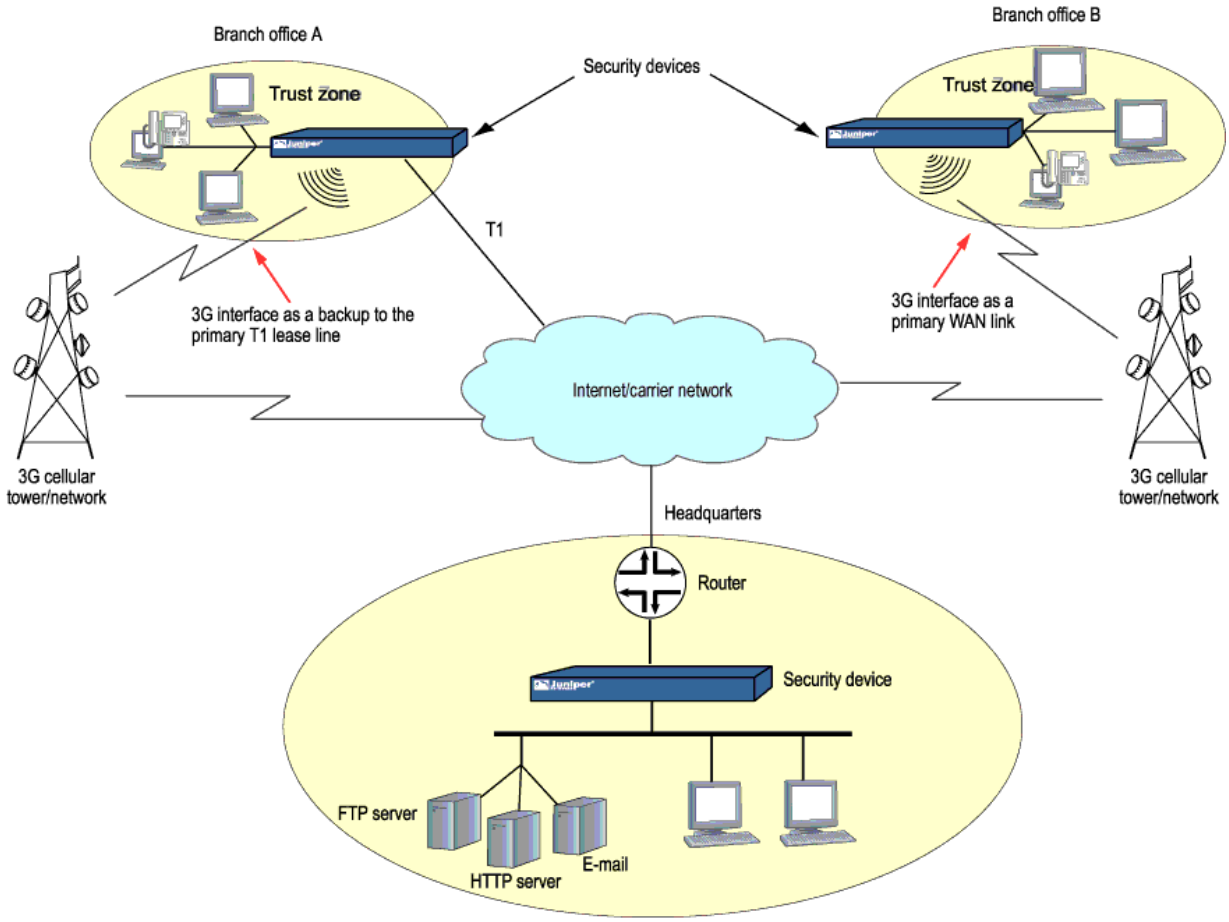

3G Wireless Modem Configuration Overview

Before you begin:

1. Install your SRX Series device and establish basic connectivity for your device. For more information, see the SRX Series Hardware Guide for your device.
2. Obtain a supported 3G wireless modem card for the device.
3. Establish an account with a cellular network service provider. Contact your service provider for more information.

4. With the services gateway powered off, insert the 3G wireless modem card into the ExpressCard slot (SRX320 devices) or 3G USB modems (SRX300 devices). Power on the device. The EXPCARD LED (for SRX320) and 3G LED (SRX320) on the front panel of the device indicates the status of the 3G wireless modem interface.

> **WARNING**: The device must be powered off before you insert the 3G wireless modem card in the ExpressCard slot (SRX320) or integrated 3G USB modem (SRX320). Do not insert or remove the card when the device is powered on.

To configure and activate the 3G wireless modem card:

1. Configure a dialer interface. See "Example: Configuring the Dialer Interface" on page 479.
2. Configure the 3G wireless modem interface. See "Example: Configuring the 3G Wireless Modem Interface" on page 488.
3. Configure security zones and policies, as needed, to allow traffic through the WAN link. See *Example: Creating Security Zones*.

To use the 3G USB modems on the SRX210 device:

1. Upgrade the BIOS software packaged inside the Junos OS image. For detailed information about BIOS upgrade procedures, see the Software Installation and Upgrade Guide.

> **NOTE**: You need the BIOS version of 2.1 or higher to use the 3G USB modems on the SRX210 device.

2. Configure the WAN port using the CLI command `set chassis routing-engine usb-wwan port 1` to enable the USB port to use the U319 USB modem.
3. Plug the 3G USB modem in to the appropriate USB slot (USB port 1) on the device.

> **NOTE**: You can use the USB modem with a standard USB extension cable of 1.8288 meters (6 ft) or longer.

4. Reboot the device to start using the 3G USB modem.

Understanding the Dialer Interface

IN THIS SECTION

- Dialer Interface Configuration Rules | 477
- Dialer Interface Authentication Support for GSM HSDPA 3G Wireless Modems | 478
- Dialer Interface Functions | 478
- Dialer Interface Operating Parameters | 479

The *dialer interface*, `dln`, is a *logical interface* for configuring properties for modem connections. You can configure multiple dialer interfaces on an SRX Series device. A dialer interface and a dialer pool (which includes the physical interface) are bound together in a dialer profile.

The dialer interface for 3G wireless modems is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

This topic contains the following sections:

Dialer Interface Configuration Rules

The following rules apply when you configure dialer interfaces for 3G wireless modem connections:

- The dialer interface must be configured to use the default Point-to-Point Protocol (PPP) encapsulation. You cannot configure Cisco High-Level Data Link Control (HDLC) or Multilink PPP (MLPPP) encapsulation on dialer interfaces.
- You cannot configure the dialer interface as a constituent link in a multilink bundle.
- You cannot configure any dial-in options for the dialer interface.

You configure the following for a dialer interface:

- A dialer pool to which the physical interface belongs.
- Source IP address for the dialer interface.
- Dial string (optional) is the destination number to be dialed.
- Authentication, for GSM HSDPA 3G wireless modem cards.
- Watch list, if the dialer interface is a backup WAN link.

With GSM HSDPA 3G wireless modem cards, you might need to configure PAP or CHAP for authentication with the service provider network. The service provider must supply the username and password, which you configure in an access profile. You then specify the access profile in a dialer interface.

Next you set the dialer interface as a backup WAN link to a primary interface. Then you create a dialer watch to enable the device to monitor the route to a head office router and set a dialer pool. Finally, you create a dialer filter firewall rule for traffic from the branch office to the main office router and associate the dialer filter with a dialer interface.

Dialer Interface Authentication Support for GSM HSDPA 3G Wireless Modems

For GSM HSDPA 3G wireless modems, you configure a dialer interface to support authentication through Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP).

CHAP is a server-driven, three-step authentication method that depends on a shared secret password that resides on both the server and the client. When you enable CHAP on a dialer interface, the device can authenticate its peer and be authenticated by its peer.

PAP allows a simple method for a peer to establish its identity using a two-way handshake during initial link establishment. After the link is established, an identification and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

Dialer Interface Functions

The dialer interface can perform backup, dialer filter, and dialer watch functions, but these operations are mutually exclusive. You can configure a single dialer interface to operate in only one of the following ways:

- As a backup interface for a single primary WAN connection. The dialer interfaces are activated only when the primary interface fails. The 3G wireless modem backup connectivity is supported on all interfaces except `lsq-0/0/0`.
- As a dialer filter. The Dialer filter enables the 3G wireless modem connection to be activated only when specific network traffic is sent on the backup WAN link. You configure a firewall rule with the dialer filter option, and then apply the dialer filter to the dialer interface.
- As a dialer watch interface. With dialer watch, the SRX Series device monitors the status of a specified route and if the route disappears, the dialer interface initiates the 3G wireless modem connection as a backup connection. To configure dialer watch, you first add the routes to be monitored to a watch list in a dialer interface; specify a dialer pool for this configuration. Then configure the 3G wireless modem interface to use the dialer pool.

Dialer Interface Operating Parameters

You can also specify optional operating parameters for the dialer interface:

- Activation delay—Number of seconds after the primary interface is down before the backup interface is activated. The default value is 0 seconds, and the maximum value is 60 seconds. Use this option only if dialer watch is configured.
- Deactivation delay—Number of seconds after the primary interface is up before the backup interface is deactivated. The default value is 0 seconds, and the maximum value is 60 seconds. Use this option only if dialer watch is configured.
- Idle timeout—Number of seconds the connection remains idle before disconnecting. The default value is 120 seconds, and the range is from 0 to 4,294,967,295 seconds.
- Initial route check—Number of seconds before the primary interface is checked to see if it is up. The default value is 120 seconds, and the range is from 1 to 300 seconds.

Example: Configuring the Dialer Interface

IN THIS SECTION

- Requirements | 479
- Overview | 480
- Configuration | 480
- Verification | 487

This example shows how to configure the dialer interface for 3G wireless modem connections.

The dialer interface for 3G wireless modems is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Requirements

Before you begin, install your SRX Series device and establish basic connectivity for your device. See "3G Wireless Modem Configuration Overview" on page 475.

Overview

In this example, you first configure the dialer interface as dl0, specify the PPP encapsulation dialer pool as 1, specify the dial string as 14691, and negotiate the address option for the interface IP address.

Configuration

IN THIS SECTION

- Configuring a Dialer Interface | **480**
- Configuring PAP on the Dialer Interface | **481**
- Configuring CHAP on the Dialer Interface | **483**
- Configuring the Dialer Interface as a Backup WAN Connection | **484**
- Configuring Dialer Watch for the 3G Wireless Modem Interface | **485**
- Configuring a Dialer Filter for the 3G Wireless Modem Interface | **486**

Configuring a Dialer Interface

CLI Quick Configuration

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

```
set interfaces dl0 description 3g-wireless encapsulation ppp unit 0 dialer-options pool 1 dial-
string 14691
set interfaces dl0 unit 0 family inet negotiate-address
```

Step-by-Step Procedure

1. Set the interface and specify the PPP encapsulation, dialer pool, and dial string.

```
[edit]
user@host# set interfaces dl0 description 3g-wireless encapsulation ppp unit 0 dialer-options
pool 1 dial-string 14691
```

2. Set the negotiate address option for the interface IP address.

```
[edit]
user@host# set interfaces dl0 unit 0 family inet negotiate-address
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces dl0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces dl0
description 3g-wireless;
encapsulation ppp;
    unit 0 {
family inet {
negotiate-address;
     }
dialer-options {
pool 1;
    dial-string 14691;
     }
}
```

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

Configuring PAP on the Dialer Interface

CLI Quick Configuration

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

```
set access profile pap-1 client clientX pap-password 7a^6b%5c
set interfaces dl0 unit 0 ppp-options pap access-profile pap-1
```

Step-by-Step Procedure

1. Configure a PAP access profile.

```
[edit]
user@host# set access profile pap-1 client clientX pap-password 7a^6b%5c
```

2. Associate the PAP access profile with a dialer interface.

```
[edit]
user@host# set interfaces dl0 unit 0 ppp-options pap access-profile pap-1
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces dl0` and `show access profile pap-1` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces dl0
unit 0 {
ppp-options {
    pap {
access-profile pap-1;
}
     }
}
[edit]
    user@host# show access profile pap-1
client clientX pap-password "$9$jnqTz3nCBESu01hSrKvZUDkqf"; ## SECRET-DATA
```

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

Configuring CHAP on the Dialer Interface

CLI Quick Configuration

With GSM HSDPA 3G wireless modem cards, you may need to configure CHAP for authentication with the service provider network. The service provider must supply the username and password, which you configure in an access profile. You then specify this access profile in a dialer interface.

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

```
set access profile chap-1 client clientX chap-secret 7a^6b%5c
set interfaces dl0 unit 0 ppp-options chap access-profile chap-1
```

Step-by-Step Procedure

1. Configure a CHAP access profile.

```
[edit]
user@host# set access profile chap-1 client clientX chap-secret 7a^6b%5c
```

2. Associate the CHAP access profile with a dialer interface.

```
[edit]
user@host# set interfaces dl0 unit 0 ppp-options chap access-profile chap-1
```

Results

From configuration mode, confirm your configuration by entering the `show access profile chap-1` and `show interfaces dl0` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show access profile chap-1
client clientX chap-secret "$9$neYpCO1REyWx-Kv87-VsYQF39Cu"; ## SECRET-DATA
[edit]
    user@host# show interfaces dl0
```

```
unit 0 {
ppp-options {
    chap {
access-profile chap-1;
}
    }
}
```

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

Configuring the Dialer Interface as a Backup WAN Connection

CLI Quick Configuration

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

```
set interfaces ge-0/0/1 unit 0 backup-options interface dl0
```

Step-by-Step Procedure

1. Set interface back up option.

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 backup-options interface dl0
```

Results

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

```
[edit]
    user@host# show interfaces ge-0/0/1
unit 0 {
backup-options {
    interface dl0.0;
```

```
    }
}
```

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

Configuring Dialer Watch for the 3G Wireless Modem Interface

CLI Quick Configuration

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

```
set interfaces dl0 description dialer-watch unit 0 dialer-options watch-list 200.200.201.1/32
set interfaces dl0 description dialer-watch unit 0 dialer-options pool dw-pool
```

Step-by-Step Procedure

1. Create a dialer watch.

```
[edit]
user@host# set interfaces dl0 description dialer-watch unit 0 dialer-options watch-list
200.200.201.1/32
```

2. Set a dialer pool.

```
[edit]
user@host# set interfaces dl0 description dialer-watch unit 0 dialer-options pool dw-pool
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces dl0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces dl0
description dialer-watch;
```

```
unit 0 {
    dialer-options {
watch-list {
200.200.201.1/32;
    }
    }
}
```

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

Configuring a Dialer Filter for the 3G Wireless Modem Interface

CLI Quick Configuration

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

```
set firewall family inet dialer-filter traffic-filter term term1 then note
```

Step-by-Step Procedure

1. Associate the dialer filter with a dialer interface.

```
[edit]
user@host# set firewall family inet dialer-filter traffic-filter term term1 then note
```

2. Check your other changes to the configuration before committing.

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

Results

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

```
[edit]
user@host# show firewall
family inet {
    dialer-filter traffic-filter {
        term term-1 {
            then note;
        }
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the Configuration | 487

Confirm that the configuration is working properly.

Verifying the Configuration

Purpose

Verify the configuration output.

Action

Verify the configuration output by entering the `show interfaces` command.

Understanding the 3G Wireless Modem Physical Interface

You configure two types of interfaces for 3G wireless modem connectivity—the physical interface and a logical dialer interface.

The physical interface for the 3G wireless modem uses the name `cl-0/0/8`. This interface is automatically created when a 3G wireless modem is installed in the device.

The 3G wireless modem physical interface is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

You configure the following properties for the physical interface:

- A dialer pool to which the physical interface belongs and the priority of the interface in the pool. A physical interface can belong to more than one dialer pool. The dialer pool priority has a range from `1` to `255`, with `1` designating the lowest-priority interfaces and `255` designating the highest-priority interfaces.
- Modem initialization string (optional). These strings begin with `AT` and execute Hayes modem commands that specify modem operation.
- GSM profile for establishing a data call with a GSM cellular network.

By default, the modem allows access to networks other than the home network.

Example: Configuring the 3G Wireless Modem Interface

IN THIS SECTION

- Requirements | **489**
- Overview | **489**
- Configuration | **489**
- Verification | **490**

This example shows how to configure the 3G wireless modem interface.

The 3G wireless modem physical interface is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Requirements

Before you begin, configure a dialer interface. See "Example: Configuring the Dialer Interface" on page 479.

Overview

In this example, you configure the physical interface as cl-0/0/8 for the 3G wireless modem to use dialer pool 1 and set the priority for the dialer pool to 25. You also configure a modem initialization string to autoanswer after two rings.

Configuration

IN THIS SECTION

- Procedure | 489

Procedure

Step-by-Step Procedure

To configure the 3G wireless modem interface:

1. Specify the dialer pool.

```
[edit]
user@host# set interfaces cl-0/0/8 dialer-options pool 1 priority 25
```

2. Specify the modem options.

```
[edit]
user@host# set interfaces cl-0/0/8 modem-options init-command-string “ATSO=2\n”
```

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces cl-0/0/8 modem options` command.

Understanding the GSM Profile

To allow data calls to a Global System for Mobile Communications (GSM) network, you must obtain the following information from your service provider:

- Username and password
- Access point name (APN)
- Whether the authentication is Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP)

You configure this information in a GSM profile associated with the 3G wireless modem physical interface. You can configure up to 16 different GSM profiles, although only one profile can be active at a time.

> **NOTE**: You also need to configure a CHAP or PAP profile with the specified username and password for the dialer interface.

Subscriber information is written to the Subscriber Identity Module (SIM) on the GSM HSDPA 3G wireless modem card. If the SIM is locked, you must unlock it before activation by using the master subsidy lock (MSL) value given by the service provider when you purchase the cellular network service.

Some service providers may preload subscriber profile information on a SIM card. The assigned subscriber information is stored in profile 1, while profile 0 is a default profile created during manufacturing. If this is the case, specify profile 1 for the GSM profile associated with the 3G wireless modem physical interface.

Configuring the information in a GSM profile associated with the 3G wireless modem physical interface is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Example: Configuring the GSM Profile

IN THIS SECTION

- Requirements | 491
- Overview | 491
- Configuration | 492
- Verification | 493

This example shows how to configure the GSM profile for the 3G wireless modem interface with service provider networks such as AT&T and T-Mobile.

NOTE: Configuring the information in a GSM profile associated with the 3G wireless modem physical interface is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Requirements

Before you begin:

- Configure a dialer interface. See "Example: Configuring the Dialer Interface" on page 479
- Configure the 3G wireless modem interface. See "Example: Configuring the 3G Wireless Modem Interface" on page 488.

Overview

IN THIS SECTION

- Topology | 492

In this example, you configure the following information provided by a service provider in a GSM profile called juniper99 that is associated with the 3G wireless modem physical interface cl-0/0/8:

- Username—`juniper99`

- Password—1@#6ahgfh
- Access point name (APN)—apn.service.com
- Authentication method—CHAP

Then you activate the profile by specifying the profile ID as profile-id 1.

Topology

Configuration

IN THIS SECTION

- Procedure | 492

Procedure

Step-by-Step Procedure

To configure a GSM profile for the 3G wireless modem interface:

1. Create a GSM profile.

```
[edit]
user@host> request modem wireless gsm create-profile profile-id 1 sip-user-id juniper99 sip-
password 16ahgfh access-point-name apn.service.com authentication-method chap
```

2. Activate the profile.

```
[edit]
user@host# set interface cl-0/0/8 cellular-options gsm-options select-profile profile-id 1
```

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

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

Verification

To verify the configuration is working properly, enter the `show interfaces cl-0/0/8` command.

Unlocking the GSM 3G Wireless Modem

Before you begin, obtain the PIN from the service provider.

The subscriber identity module (SIM) in the GSM 3G wireless modem card is a detachable smart card. Swapping out the SIM allows you to change the service provider network, however some service providers lock the SIM to prevent unauthorized access to the service provider's network. If this is the case, you will need to unlock the SIM by using an personal identification number (PIN), a four-digit number provided by the service provider.

NOTE: Unlocking the SIM in a 3G wireless modem card is not supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Use the CLI operational mode command to unlock the SIM on the GSM 3G wireless modem card.

This example uses the PIN `3210` from the service provider.

To unlock the SIM on the GSM 3G wireless modem card:

```
user@host> request modem wireless gsm sim-unlock cl-0/0/8 pin 3210
```

A SIM is blocked after three consecutive failed unlock attempts; this is a security feature to prevent brute force attempts to unlock the SIM. When the SIM is blocked, you need to unblock the SIM with an eight-digit PIN unlocking key (PUK) obtained from the service provider.

To unlock the SIM automatically on reboot:

```
user@host# set interfaces cl-0/0/8 cellular-options gsm-options sim-unlock-code
Enter PIN:
user@host#
```

NOTE: On SRX300, SRX320 devices, when you power on or reboot the device, the Subscriber Identity Module (SIM) will be locked. If the SIM Personal Identification Number (PIN) or the

unlock code is configured in the `set interfaces cl-0/0/8 cellular-options gsm-options sim-unlock-code` configuration command, then Junos OS attempts to unlock the SIM only once. This is to keep the SIM from being blocked. If the SIM is blocked, you must provide a PIN Unblocking Key (PUK) obtained from the service provider. If the wrong SIM PIN is configured, the SIM will remain locked, and the administrator can unlock it by using the remaining two attempts.

Use the CLI operational mode command to unblock the SIM.

This example uses the PUK `76543210` from the service provider.

To unblock the SIM:

```
user@host> request modem wireless gsm sim-unblock cl-0/0/8 puk 76543210
```

NOTE: If you enter the PUK incorrectly ten times, you will need to return the SIM to the service provider for reactivation.

Configuring CDMA EV-DO Modem Cards

IN THIS SECTION

- Understanding Account Activation for CDMA EV-DO Modem Cards | **495**
- Activating the CDMA EV-DO Modem Card Manually | **497**
- Activating the CDMA EV-DO Modem Card with IOTA Provisioning | **499**
- Activating the CDMA EV-DO Modem Card with OTASP Provisioning | **499**

The below topics discuss the account activation for CDMA EV-DO Modem Cards and activation details on security devices.

Understanding Account Activation for CDMA EV-DO Modem Cards

IN THIS SECTION

- Obtaining Electronic Serial Number (ESN) | 495
- Account Activation Modes | 496

Account activation is the process of enabling the CDMA EV-DO wireless modem card to connect to your service provider's cellular network. This is a one-time process where your subscriber information is saved in nonvolatile memory on the card. The procedure you use to perform account activation depends upon the service provider network.

NOTE: Activating an account for a CDMA EV-DO 3G wireless modem card is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Before activating an account, you can verify the signal strength on the 3G wireless modem interface by using the `show modem wireless interface cl-0/0/8 rssi` command. The signal strength should be at least -90 dB and preferably better than -80 dB (-125 dB indicates nil signal strength). If the signal strength is below -90 dB, activation may not be possible from that location. For example:

```
user@host> show modem wireless interface cl-0/0/8 rssi
Current Radio Signal Strength (RSSI) = -98 dBm
```

This topic contains the following sections:

Obtaining Electronic Serial Number (ESN)

The service provider requires the electronic serial number (ESN) of the 3G wireless modem card to activate your account and to generate the necessary information you need to activate the card. You can obtain the ESN number of the modem card in the following ways:

- Inspect the modem card itself; the ESN is printed on the card.

- Use the CLI `show modem wireless interface cl-0/0/8 firmware` command, as shown in the following example, and note the value for the Electronic Serial Number (ESN) field:

```
user@host> show modem wireless interface cl-0/0/8 firmware
Modem Firmware Version : p2005600
Modem Firmware built date : 12-09-07
Card type : Aircard 597E - CDMA EV-DO revA
Manufacturer : Sierra Wireless, Inc.
Hardware Version : 1.0
Electronic Serial Number (ESN) : 0x6032688F
Preferred Roaming List (PRL) Version : 20224
Supported Mode : 1xev-do rev-a, 1x
Current Modem Temperature : 32 degrees Celsius
Modem Activated : YES
Activation Date: 2-06-08
Modem PIN Security : Unlocked
Power-up lock : Disabled
```

Account Activation Modes

For the CDMA EV-DO 3G wireless modem card, account activation can be done through one or more of the following modes:

- Over the air service provisioning (OTASP)—protocol for programming phones over the air using Interim Standard 95 (IS-95) Data Burst Messages.

  To activate the 3G wireless modem card with OTASP, you need to obtain from the service provider the dial number that the modem will use to contact the network. Typically, OTASP dial numbers begin with the feature code *228 to indicate an activation call type to the cellular network's base transceiver station, followed by additional digits specified by the service provider.

- Internet-based over the air (IOTA) provisioning—method for programming phones for voice and data services

- Manually providing the required information by entering in a CLI *operational mode command*

Sprint uses manual and IOTA activation, whereas Verizon uses only OTASP.

NOTE: The 3G wireless modem is set into Single-Carrier Radio Transmission Technology (1xRTT) mode automatically when it is activated for Verizon networks.

Activating the CDMA EV-DO Modem Card Manually

Before you begin, the service provider must activate your account before you can activate the CDMA EV-DO 3G wireless modem card.

Manual activation stores the supplied values into the 3G wireless modem card's nonvolatile memory. This topic describes the activation of the CDMA EV-DO 3G wireless modem card for use with service provider networks such as Sprint.

NOTE: Activating a CDMA EV-DO 3G wireless modem card manually is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Using the electronic serial number (ESN) you provided and your account information, the service provider supplies you with the following information for manual activation of the 3G wireless modem card:

- Master subsidy lock (MSL)—activation code
- Mobile directory number (MDN)—10-digit user phone number
- International mobile station identify (IMSI)—Mobile subscriber information
- Simple IP user identification (SIP-ID)—Username
- Simple IP password (SIP-Password)—Password

You also need to obtain the following information from the 3G wireless modem card itself for the activation:

- System identification (SID)—Number between 0 and 32767
- Network identification (NID)—Number between 0 and 65535

Use the CLI `show modem wireless interface cl-0/0/8 network` command to display the SID and NID, as shown in the following example:

```
user@host> show modem wireless interface cl-0/0/8 network
Running Operating mode : 1xEV-DO (Rev A) and 1xRTT
Call Setup Mode : Mobile IP only
System Identifier (SID) : 3421
Network Identifier (NID) : 91
Roaming Status(1xRTT) : Home
```

```
Idle Digital Mode : HDR
System Time : Wed Jun6 15:16:9 2008
```

Use the CLI operational mode command to manually activate the 3G wireless modem card.

This example uses the following values for manual activation:

- MSL (from service provider)—43210
- MDN (from service provider)—0123456789
- IMSI (from service provider)—0123456789
- SIP-ID (from service provider)—jnpr
- SIP-Password (from service provider)—jn9rl
- SID (from modem card)—12345
- NID (from modem card)—12345

To activate the CDMA EV-DO 3G wireless modem card manually:

```
user@host> request modem wireless interface cl-0/0/8 activate manual msl 43210 mdn 0123456789
imsi 0123456789 sid 12345 nid 12345 sip-id jnpr sip-password jn9rl
Checking status...
Modem current activation status: Not Activated
Starting activation...
Performing account activation step 1/6 : [Unlock] Done
Performing account activation step 2/6 : [Set MDN] Done
Performing account activation step 3/6 : [Set SIP Info] Done
Performing account activation step 4/6 : [Set IMSI] Done
Performing account activation step 5/6 : [Set SID/NID] Done
Performing account activation step 6/6 : [Commit/Lock] Done
Configuration Commit Result: PASS
Resetting the modem ... Done
Account activation in progress. It can take up to 5 minutes
Please check the trace logs for details.
```

To check the trace log for account activation details:

```
user@host> tail -f /var/log/wwand.log
Jun 25 04:42:55: IOTA cl-0/0/8 Event: IOTA Start... Success
Jun 25 04:43:45: IOTA cl-0/0/8 OTA SPL unlock...  Success
```

```
Jun 25 04:43:56: IOTA cl-0/0/8 Committing OTA Parameters to NVRAM... Success
Jun 25 04:44:02: IOTA cl-0/0/8 Over the air provisioning... Complete
Jun 25 04:44:04: IOTA cl-0/0/8 IOTA Event: IOTA End... Success
```

Activating the CDMA EV-DO Modem Card with IOTA Provisioning

Before you begin, activate the CDMA EV-DO 3G wireless modem card. See "Understanding Account Activation for CDMA EV-DO Modem Cards" on page 495.

Manual activation stores the supplied values in the 3G wireless modem card's nonvolatile memory. If the modem card is reset or you need to update Mobile IP (MIP) parameters, use the CLI operational mode command to activate the modem card with IOTA.

NOTE: Activating a CDMA EV-DO 3G wireless modem card with IOTA provisioning is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

To activate the CDMA EV-DO 3G wireless modem card with IOTA:

```
user@host> request modem wireless interface cl-0/0/8 activate iota
Beginning IOTA Activation. It can take up to 5 minutes
Please check the trace logs for details.
```

To check the trace log for account activation details:

```
user@host> tail -f /var/log/wwand.log
Jun 25 04:42:55: IOTA cl-0/0/8 Event: IOTA Start... Success
Jun 25 04:43:45: IOTA cl-0/0/8 OTA SPL unlock...  Success
Jun 25 04:43:56: IOTA cl-0/0/8 Committing OTA Parameters to NVRAM... Success
Jun 25 04:44:02: IOTA cl-0/0/8 Over the air provisioning... Complete
Jun 25 04:44:04: IOTA cl-0/0/8 IOTA Event: IOTA End... Success
```

Activating the CDMA EV-DO Modem Card with OTASP Provisioning

Before you begin:

- Obtain the dial number that the modem will use to contact the network from the service provider.
- The service provider must activate your account before OTASP provisioning can proceed.

This topic describes the activation of the CDMA EV-DO 3G wireless modem card for use with service provider networks such as Verizon.

> **NOTE**: Activating a CDMA EV-DO 3G wireless modem card with OTASP provisioning is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Use the CLI operational mode command to activate the 3G wireless modem card.

In this example, the dial number from the service provider is `*22864`.

To activate the CDMA EV-DO 3G wireless modem card with OTASP provisioning:

```
user@host> request modem wireless interface cl-0/0/8 activate otasp dial-string *22864
OTASP number *22286*, Selecting NAM 0
Beginning OTASP Activation. It can take up to 5 minutes
Please check the trace logs for details.
```

To check the trace log for account activation details:

```
user@host> tail -f /var/log/wwand.log
Jun 25 04:42:55: OTASP cl-0/0/8 OTA SPL unlock... Success
Jun 25 04:43:42: OTASP cl-0/0/8 OTA PRL download... Success
Jun 25 04:43:55: OTASP cl-0/0/8 OTA Profile downloaded... Success
Jun 25 04:43:58: OTASP cl-0/0/8 OTA MDN download... Success
Jun 25 04:44:04: OTASP cl-0/0/8 Committing OTA Parameters to NVRAM... Success
Jun 25 04:44:45: Over the air provisioning... Complete
```

Configuring USB Modems for Dial Backup

IN THIS SECTION

- USB Modem Interface Overview | 501
- USB Modem Configuration Overview | 505
- Example: Configuring a USB Modem Interface | 507
- Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup | 511
- Example: Configuring a Dialer Interface for USB Modem Dial-In | 521
- Example: Configuring PAP on Dialer Interfaces | 524
- Example: Configuring CHAP on Dialer Interfaces | 525

The topics below discuss the USB modem interfaces, its configuration details, examples of configuring dialer interface, configuring PAP on dialer interface and CHAP on dialer interface.

USB Modem Interface Overview

IN THIS SECTION

- USB Modem Interfaces | 502
- Dialer Interface Rules | 503
- How the Device Initializes USB Modems | 503

Juniper Networks SRX Series devices support the use of USB modems for remote management. You can use Telnet or SSH to connect to the device from a remote location through two modems over a telephone network. The USB modem is connected to the USB port on the device, and a second modem is connected to a remote management device such as a PC or laptop computer.

NOTE: USB modems are no longer supported for dial backup on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

You can configure your device to fail over to a USB modem connection when the primary Internet connection experiences interruption.

A USB modem connects to a device through modem interfaces that you configure. The device applies its own modem AT commands to initialize the attached modem. Modem setup requires that you connect and configure the USB modem at the device and the modem at the user end of the network.

You use either the J-Web configuration editor or CLI configuration editor to configure the USB modem and its supporting dialer interfaces.

NOTE: Low-latency traffic such as VoIP traffic is not supported over USB modem connections.

NOTE: We recommend using a US Robotics USB 56k V.92 Modem, model number USR Model 5637.

USB Modem Interfaces

You configure two types of interfaces for USB modem connectivity:

- A physical interface which uses the naming convention `umd0`. The device creates this interface when a USB modem is connected to the USB port.
- A *logical interface* called the dialer interface. You use the dialer interface, `dl`*n*, to configure dialing properties for USB modem connections. The dialer interface can be configured using Point-to-Point Protocol (PPP) encapsulation. You can also configure the dialer interface to support authentication protocols—PPP Challenge Handshake (CHAP) or Password Authentication Protocol (PAP). You can configure multiple dialer interfaces for different functions on the device. After configuring the dialer interface, you must configure a backup method such as a dialer backup, a dialer filter, or a dialer watch.

The USB modem provides a dial-in remote management interface, and supports dialer interface features by sharing the same dial pool as a dialer interface. The dial pool allows the logical dialer interface and the physical interface to be bound together dynamically on a per-call basis. You can configure the USB modem to operate either as a dial-in console for management or as a dial-in WAN backup interface.

Dialer pool priority has a range from 1 to 255, with 1 designating the lowest priority interfaces and 255 designating the highest priority interfaces.

Dialer Interface Rules

The following rules apply when you configure dialer interfaces for USB modem connections:

- The dialer interface must be configured to use PPP encapsulation. You cannot configure Cisco High-Level Data Link Control (HDLC) or Multilink PPP (MLPPP) encapsulation on dialer interfaces.
- The dialer interface cannot be configured as a constituent link in a multilink bundle.
- The dialer interface can perform backup, dialer filter, and dialer watch functions, but these operations are mutually exclusive. You can configure a single dialer interface to operate in only one of the following ways:
  - As a backup interface—for one primary interface
  - As a dialer filter
  - As a dialer watch interface

The backup dialer interfaces are activated only when the primary interface fails. USB modem backup connectivity is supported on all interfaces except lsq-0/0/0.

The dial-on-demand routing backup method allows a USB modem connection to be activated only when network traffic configured as an “interesting packet” arrives on the network. Once the network traffic is sent, an inactivity timer is triggered and the connection is closed. You define an interesting packet using the dialer filter feature of the device. To configure dial-on-demand routing backup using a dialer filter, you first configure the dialer filter and then apply the filter to the dialer interface.

Dialer watch is a backup method that integrates backup dialing with routing capabilities and provides reliable connectivity without relying on a dialer filter to trigger outgoing USB modem connections. With dialer watch, the device monitors the existence of a specified route. If the route disappears, the dialer interface initiates the USB modem connection as a backup connection.

How the Device Initializes USB Modems

When you connect the USB modem to the USB port on the device, the device applies the modem AT commands configured in the `init-command-string` command to the initialization commands on the modem.

If you do not configure modem AT commands for the `init-command-string` command, the device applies the following default sequence of initialization commands to the modem: `AT S7=45 S0=0 V1 X4 &C1 E0 Q0 &Q8 %C0`. Table 46 on page 504 describes the commands. For more information about these commands, see the documentation for your modem.

Table 46: Default Modem Initialization Commands

Modem Command	Description
`AT`	Attention. Informs the modem that a command follows.
`S7=45`	Instructs the modem to wait 45 seconds for a telecommunications service provider (carrier) signal before terminating the call.
`S0=0`	Disables the auto answer feature, whereby the modem automatically answers calls.
`V1`	Displays result codes as words.
`&C1`	Disables reset of the modem when it loses the carrier signal.
`E0`	Disables the display on the local terminal of commands issued to the modem from the local terminal.
`Q0`	Enables the display of result codes.
`&Q8`	Enables Microcom Networking Protocol (MNP) error control mode.
`%C0`	Disables data compression.

When the device applies the modem AT commands in the `init-command-string` command or the default sequence of initialization commands to the modem, it compares them to the initialization commands already configured on the modem and makes the following changes:

- If the commands are the same, the device overrides existing modem values that do not match. For example, if the initialization commands on the modem include `S0=0` and the device's `init-command-string` command includes `S0=2`, the device applies `S0=2`.
- If the initialization commands on the modem do not include a command in the device's `init-command-string` command, the device adds it. For example, if the `init-command-string` command includes the command `L2`, but the modem commands do not include it, the device adds `L2` to the initialization commands configured on the modem.

NOTE: On SRX210 devices, the USB modem interface can handle bidirectional traffic of up to 19 Kbps. On oversubscription of this amount (that is, bidirectional traffic of 20 Kbps or above), keepalives do not get exchanged, and the interface goes down. (Platform support depends on the Junos OS release in your installation.)

USB Modem Configuration Overview

NOTE: USB modems are no longer supported for dial backup on SRX300, SRX320, SRX340, and SRX345 devices.

Before you begin:

1. Install device hardware. For more information, see the Getting Started Guide for your device.
2. Establish basic connectivity. For more information, see the Getting Started Guide for your device.
3. Order a US Robotics USB 56k V.92 Modem, model number USR Model 5637 (http://www.usr.com/).
4. Order a public switched telephone network (PSTN) line from your telecommunications service provider. Contact your service provider for more information.
5. Connect the USB modem to the device's USB port.

   NOTE: When you connect the USB modem to the USB port on the device, the USB modem is initialized with the modem initialization string configured for the USB modem interface on the device.

   a. Plug the modem into the USB port.

   b. Connect the modem to your telephone network.

      i.

Suppose you have a branch office router and a head office router each with a USB modem interface and a dialer interface. This example shows you how to establish a backup connection between the branch office and head office routers. See Table 47 on page 506 for a summarized description of the procedure.

Table 47: Configuring Branch Office and Head Office Routers for USB Modem Backup Connectivity

Router Location	Configuration Requirement	Procedure
Branch Office	Configure the logical dialer interface on the branch office router for USB modem dial backup.	To configure the logical dialer interface, see *Example: Configuring a USB Modem Interface*.
	Configure the dialer interface `dl0` on the branch office router using one of the following backup methods: • Configure the dialer interface `dl0` as the backup interface on the branch office router's primary T1 interface `t1-1/0/0`. • Configure a dialer filter on the branch office router's dialer interface. • Configure a dialer watch on the branch office router's dialer interface.	Configure the dialer interface using one of the following backup methods: • To configure `dl0` as a backup for `t1-1/0/0` see "Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup" on page 511. • To configure a dialer filter on `dl0`, see "Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup" on page 511. • To configure a dialer watch on `dl0`, see "Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup" on page 511.
Head Office	Configure dial-in on the dialer interface `dl0` on the head office router.	To configure dial-in on the head office router, see *Example: Configuring a Dialer Interface for USB Modem Dial-In*.

If the dialer interface is configured to accept only calls from a specific caller ID, the device matches the incoming call's caller ID against the caller IDs configured on its dialer interfaces. If an exact match is not found and the incoming call's caller ID has more digits than the configured caller IDs, the device performs a right-to-left match of the incoming call's caller ID with the configured caller IDs and accepts the incoming call if a match is found. For example, if the incoming call's caller ID is 4085321091 and the caller ID configured on a dialer interface is 5321091, the incoming call is accepted. Each dialer interface accepts calls from only callers whose caller IDs are configured on it.

See Table 48 on page 507 for a list of available incoming map options.

Table 48: Incoming Map Options

Option	Description
accept-all	Dialer interface accepts all incoming calls. You can configure the `accept-all` option for only one of the dialer interfaces associated with a USB modem physical interface. The dialer interface with the `accept-all` option configured is used only if the incoming call's caller ID does not match the caller IDs configured on other dialer interfaces.
caller	Dialer interface accepts calls from a specific caller ID. You can configure a maximum of 15 caller IDs per dialer interface. The same caller ID must not be configured on different dialer interfaces. However, you can configure caller IDs with more or fewer digits on different dialer interfaces. For example, you can configure the caller IDs 14085551515, 4085551515, and 5551515 on different dialer interfaces.

You configure dialer interfaces to support PAP. PAP allows a simple method for a peer to establish its identity using a two-way handshake during initial link establishment. After the link is established, an ID and password pair are repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

Example: Configuring a USB Modem Interface

IN THIS SECTION

- Requirements | 508
- Overview | 508
- Configuration | 508
- Verification | 510

This example shows how to configure a USB modem interface for dial backup.

NOTE: USB modems are no longer supported for dial backup on SRX300, SRX320, SRX340, and SRX345 devices.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you create an interface called as umd0 for USB modem connectivity and set the dialer pool priority to 25. You also configure a modem initialization string to autoanswer after a specified number of rings. The default modem initialization string is `AT S7=45 S0=0 V1 X4 &C1 E0 Q0 &Q8 %C0`. The modem command `S0=0` disables the modem from autoanswering the calls. Finally, you set the modem to act as a dial-in WAN backup interface.

Configuration

IN THIS SECTION

- Procedure | 508

Procedure

CLI Quick Configuration

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

```
set interfaces umd0 dialer-options pool usb-modem-dialer-pool priority 25
set modem-options init-command-string "ATS0=2 \n" dialin routable
```

Step-by-Step Procedure

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

To configure a USB modem interface for dial backup:

1. Create an interface.

```
[edit]
user@host# edit interfaces umd0
```

2. Set the dialer options and priority.

```
[edit interfaces umd0]
user@host# set dialer-options pool usb-modem-dialer-pool priority 25
```

3. Specify the modem options.

```
[edit interfaces umd0]
user@host# set modem-options init-command-string "ATS0=2 \n"
```

4. Set the modem to act as a dial-in WAN backup interface.

```
[edit interfaces umd0]
user@host# set modem-options dialin routable
```

Results

From configuration mode, confirm your configuration by entering the `show interface umd0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interface umd0
modem-options {
init-command-string "ATS0=2 \n";
    dialin routable;
```

```
}
dialer-options {
    pool usb-modem-dialer-pool priority 25;
}
```

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

Verification

IN THIS SECTION

- Verifying the Configuration | **510**

Confirm that the configuration is working properly.

Verifying the Configuration

Purpose

Verify a USB modem interface for dial backup.

Action

From configuration mode, enter the `show interfaces umd0 extensive` command. The output shows a summary of interface information and displays the modem status.

```
Physical interface:   umd0, Enabled, Physical link is Up
Interface index:      64, SNMP ifIndex: 33, Generation: 1
  Type: Async-Serial, Link-level type: PPP-Subordinate, MTU: 1504,
Clocking: Unspecified, Speed: MODEM
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  Last flapped   : Never
  Statistics last cleared: Never
 Traffic statistics:
   Input  bytes  :               21672
```

```
    Output bytes  :                22558
    Input  packets:                 1782
    Output packets:                 1832
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0,
Resource errors: 0
  Output errors:
    Carrier transitions: 63, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  MODEM status:
    Modem type                    : LT V.92 1.0 MT5634ZBA-USB-V92 Data/Fax Modem
(Dual Config) Version 2.27m
    Initialization command string : ATS0=2
    Initialization status         : Ok
    Call status                   : Connected to 4085551515
    Call duration                 : 13429 seconds
    Call direction                : Dialin
    Baud rate                     : 33600 bps
    Most recent error code        : NO CARRIER

  Logical interface umd0.0 (Index 2) (SNMP ifIndex 34) (Generation 1)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP-Subordinate
```

Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup

IN THIS SECTION

- Requirements | **512**
- Overview | **512**
- Configuration | **512**
- Verification | **521**

This example shows how to configure a dialer interfaces and backup methods for USB modem dial backup.

NOTE: USB modems are no longer supported for dial backup on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Requirements

Before you begin, configure a USB modem for the device. See *Example: Configuring a USB Modem Interface*.

Overview

IN THIS SECTION

- Topology | **512**

In this example, you configure a logical dialer interface on the branch office router for the USB modem dial backup. You then configure dial backup to allow one or more dialer interfaces to be configured as the backup link for the primary serial interface. To configure dialer watch, you first add a dialer watch interface and then configure the USB modem interface to participate as a dialer watch interface. The USB modem interface must have the same pool identifier to participate in dialer watch. Dialer pool name dw-pool is used when configuring the USB modem interface.

Topology

Configuration

IN THIS SECTION

- Configuring a Dialer Interface for USB Modem Dial Backup | **513**
- Configuring a Dial Backup for a USB Modem Connection | **515**
- Configuring a Dialer Filter for USB Modem Dial Backup | **517**
- Configuring a Dialer Watch for USB Modem Dial Backup | **519**

Configuring a Dialer Interface for USB Modem Dial Backup

CLI Quick Configuration

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

```
set interfaces dl0 description USB-modem-backup encapsulation ppp
set interfaces dl0 unit 0 dialer-options activation-delay 60 deactivation-delay 30 idle-timeout
30 initial-route-check 30 pool usb-modem-dialer-pool
set interfaces dl0 unit 0 dialer-options dial-string 5551212
set interfaces dl0 unit 0 family inet address 172.20.10.2 destination 172.20.10.1
```

Step-by-Step Procedure

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

To configure a logical dialer interface on the branch office router for the USB modem dial backup:

1. Create an interface.

```
[edit]
user@host# edit interfaces dl0
```

2. Specify a description.

```
[edit interfaces dl0]
user@host# set description USB-modem-backup
```

3. Configure PPP encapsulation.

```
[edit interfaces dl0]
user@host# set encapsulation ppp
```

NOTE: You cannot configure Cisco High-Level Data Link Control (HDLC) or Multilink PPP (MLPPP) encapsulation on dialer interfaces used in USB modem connections.

4. Create the logical unit.

```
[edit interfaces dl0]
user@host# set unit 0
```

NOTE: You can set the logical unit to 0 only.

5. Configure the dialer options.

```
[edit interfaces dl0]
user@host# edit unit 0 dialer-options
user@host# set activation-delay 60
user@host# set deactivation-delay 30
user@host# set idle-timeout 30 initial-route-check 30 pool usb-modem-dialer-pool
```

6. Configure the telephone number of the remote destination.

```
[edit interfaces dl0 unit 0 dialer-options]
user@host# set dial-string 5551212
```

7. Configure source and destination IP addresses.

```
[edit]
user@host# edit interfaces dl0 unit 0
user@host# set family inet address 172.20.10.2 destination 172.20.10.1
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces dl0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces dl0
    description USB-modem-backup;
encapsulation ppp;
unit 0 {
family inet {
    address 172.20.10.2/32 {
        destination 172.20.10.1;
        }
    }
        dialer-options {
            pool usb-modem-dialer-pool;
        dial-string 5551212;
            idle-timeout 30;
                activation-delay 60;
                deactivation-delay 30;
                initial-route-check 30;
                }
            }
```

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

Configuring a Dial Backup for a USB Modem Connection

CLI Quick Configuration

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

```
set interfaces t1-1/0/0 unit 0 backup-options interface dl0.0
```

Step-by-Step Procedure

To configure a dial backup for a USB modem connection:

1. Select the physical interface.

```
[edit]
user@host# edit interfaces t1-1/0/0 unit 0
```

2. Configure the backup dialer interface.

```
[edit]
user@host# set backup-options interface dl0.0
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces t1-1/0/0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host#  show interfaces t1-1/0/0
    encapsulation ppp;
    unit 0 {
        backup-options {
    interface dl0.0;
    }
}
```

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

Configuring a Dialer Filter for USB Modem Dial Backup

CLI Quick Configuration

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

```
set firewall family inet dialer-filter interesting-traffic term term1 from source-address
20.20.90.4/32
set firewall family inet dialer-filter interesting-traffic term term1 from destination-address
200.200.201.1/32
set firewall family inet dialer-filter interesting-traffic term term1 then note
set interfaces dl0 unit 0 family inet filter dialer interesting-traffic
```

Step-by-Step Procedure

To configure a dialer filter for USB modem dial backup:

1. Create an interface.

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

2. Configure the dialer filter name.

```
[edit]
user@host# edit family inet
user@host# edit dialer-filter interesting-traffic
```

3. Configure the dialer filter rule name and term behavior.

```
[edit]
user@host# edit term term1
user@host# set from source-address 20.20.90.4/32
user@host# set from destination-address 200.200.201.1/32
```

4. Configure the then part of the dialer filter.

```
[edit]
user@host# set then note
```

5. Select the dialer interface to apply the filter.

```
[edit]
user@host# edit interfaces dl0 unit 0
```

6. Apply the dialer filter to the dialer interface.

```
[edit]
user@host# edit family inet filter
user@host# set dialer interesting-traffic
```

Results

From configuration mode, confirm your configuration by entering the `show firewall family inet dialer-filter interesting-traffic` and `show interfaces dl0`commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show firewall family inet dialer-filter interesting-traffic
    term term1 {
from {
    source-address {
        20.20.90.4/32;
    }
        destination-address {
            200.200.201.1/32;
        }
        }
            then note;
        }
[edit]
    user@host# show interfaces dl0
unit 0 {
```

```
family inet {
    filter {
dialer interesting-traffic;
    }
        }
    }
```

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

Configuring a Dialer Watch for USB Modem Dial Backup

CLI Quick Configuration

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

```
set interfaces dl0 description dialer-watch unit 0 dialer-options watch-list 200.200.201.1/32
set interfaces dl0 unit 0 dialer-options pool dw-pool
set interfaces umd0 dialer-options pool dw-pool
```

Step-by-Step Procedure

To configure a dialer watch for USB modem dial backup:

1. Create an interface.

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

2. Specify a description.

```
[edit]
user@host# edit dl0
user@host# set description dialer-watch
```

3. Configure the route to the head office router for dialer watch.

```
[edit]
user@host# edit unit 0 dialer-options
user@host# set watch-list 200.200.201.1/32
```

4. Configure the name of the dialer pool.

```
[edit]
user@host# set pool dw-pool
```

5. Select the USB modem physical interface.

```
[edit]
user@host# edit interfaces umd0 dialer-options pool dw-pool
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces dl0` and `show interfaces umd0` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces dl0
    dialer-options {
        pool dw-pool;
    }
[edit]
user@host# show interfaces umd0
    description dialer-watch;
        unit 0 {
    dialer-options {
        pool dw-pool;
            watch-list {
        200.200.201.1/32;
    }
        }
    }
```

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

Verification

IN THIS SECTION

- Verifying the Configuration | 521

Confirm that the configuration is working properly.

Verifying the Configuration

Purpose

Verify the configuration output.

Action

From operational mode, enter the `show interface terse` command.

Example: Configuring a Dialer Interface for USB Modem Dial-In

IN THIS SECTION

- Requirements | 522
- Overview | 522
- Configuration | 523
- Verification | 524

This example shows how to configure a dialer interface for USB modem dial-in.

NOTE: USB modems are no longer supported for dial-in to a dialer interface on SRX300, SRX320, SRX340, and SRX345 devices.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

To enable connections to the USB modem from a remote location, you must configure the dialer interfaces set up for USB modem use to accept incoming calls. You can configure a dialer interface to accept all incoming calls or accept only calls from one or more caller IDs.

If the dialer interface is configured to accept only calls from a specific caller ID, the system matches the incoming call's caller ID against the caller IDs configured on its dialer interfaces. If an exact match is not found and the incoming call's caller ID has more digits than the configured caller IDs, the system performs a right-to-left match of the incoming call's caller ID with the configured caller IDs and accepts the incoming call if a match is found. For example, if the incoming call's caller ID is 4085550115 and the caller ID configured on a dialer interface is 5550115, the incoming call is accepted. Each dialer interface accepts calls from only callers whose caller IDs are configured on it.

You can configure the following incoming map options for the dialer interface:

- `accept-all`—Dialer interface accepts all incoming calls.

  You can configure the `accept-all` option for only one of the dialer interfaces associated with a USB modem physical interface. The device uses the dialer interface with the `accept-all` option configured only if the incoming call's caller ID does not match the caller IDs configured on other dialer interfaces.

- `caller`—Dialer interface accepts calls from a specific caller ID— for example, `4085550115`. You can configure a maximum of 15 caller IDs per dialer interface.

  The same caller ID must not be configured on different dialer interfaces. However, you can configure caller IDs with more or fewer digits on different dialer interfaces. For example, you can configure the caller IDs 14085550115, 4085550115, and 5550115 on different dialer interfaces.

In this example, you configure the incoming map option as caller 4085550115 for dialer interface dl0.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 523
- Procedure | 523

CLI Quick Configuration

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

```
set interfaces dl0 unit 0 dialer-options incoming-map caller 4085550115
```

Procedure

Step-by-Step Procedure

To configure a dialer interface for USB modem dial-in:

1. Select a dialer interface.

```
[edit]
user@host# edit interfaces dl0
```

2. Configure the incoming map options.

```
[edit]
user@host# edit unit 0 dialer-options incoming-map caller 4085551515
```

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

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

Verification

To verify the configuration is working properly, enter the `show interface dl0` command.

Example: Configuring PAP on Dialer Interfaces

IN THIS SECTION

- Requirements | 524
- Overview | 524
- Configuration | 524
- Verification | 525

This example shows how to configure PAP on dialer interfaces.

NOTE: Configuring PAP on dialer interfaces is no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you specify a PAP access profile with a client username and a PAP password and select a dialer interface. Finally, you configure PAP on the dialer interface and specify the local name and password.

Configuration

IN THIS SECTION

- Procedure | 525

Procedure

Step-by-Step Procedure

To configure PAP on the dialer interface:

1. Specify a PAP access profile.

```
[edit]
user@host# set access profile pap-access-profile client pap-access-user pap-password my-pap
```

2. Select a dialer interface.

```
[edit]
user@host# edit interfaces dl0 unit 0
```

3. Configure PAP on the dialer interface.

```
[edit]
user@host# set ppp-options pap local-name pap-access-user local-password my-pap
```

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

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

Verification

To verify the configuration is working properly, enter the `show interface dl0` command.

Example: Configuring CHAP on Dialer Interfaces

IN THIS SECTION

Requirements | 526

- Overview | 526
- Configuration | 526
- Verification | 527

This example shows how to configure CHAP on dialer interfaces for authentication.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you configure dialer interfaces to support CHAP for authentication. CHAP is a server-driven, three-step authentication method that depends on a shared secret password residing on both the server and the client. You specify a CHAP access profile with a client username and a password. You then specify a dialer interface as dl0. Finally, you enable CHAP on a dialer interface and specify a unique profile name containing a client list and access parameters.

Configuration

IN THIS SECTION

- Procedure | 526

Procedure

Step-by-Step Procedure

To configure CHAP on a dialer interface:

1. Specify a CHAP access profile.

```
[edit]
user@host# set access profile usb-modem-access-profile client usb-modem-user chap-secret my-secret
```

2. Select a dialer interface.

```
[edit]
user@host# edit interfaces dl0 unit 0
```

3. Enable CHAP on the dialer interface.

```
[edit]
user@host# set ppp-options chap access-profile usb-modem-access-profile
```

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

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

Verification

To verify the configuration is working properly, enter the `show interface dl0` command.

Configuring DOCSIS Mini-PIM Interfaces

IN THIS SECTION

- DOCSIS Mini-PIM Interface Overview | **528**
- Software Features Supported on DOCSIS Mini-PIMs | **530**
- Example: Configuring the DOCSIS Mini-PIM Interfaces | **531**

Data over Cable Service Interface Specifications (DOCSIS) define the communications and operation support interface requirements for a data-over-cable system. The topics below discuss the overview of DOCSIS Mini-PIM interface, its configuration details, and software features supported on DOCSIS Mini-PIM interfaces on SRX series devices.

DOCSIS Mini-PIM Interface Overview

Data over Cable Service Interface Specifications (DOCSIS) define the communications and operation support interface requirements for a data-over-cable system. Cable operators use DOCSIS to provide Internet access over their existing cable infrastructure for both residential and business customers. DOCSIS 3.0 is the latest interface standard, allowing channel bonding to deliver speeds higher than 100 Mbps throughput in either direction, far surpassing other WAN technologies such as T1/E1, ADSL2+, ISDN, and DS3.

NOTE: On SRX210 Services Gateway, the DOCSIS Mini-PIM delivers speeds up to a maximum of 100 Mbps throughput in each direction.

NOTE: DOCSIS Mini-PIM interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

DOCSIS network architecture includes a cable modem on SRX Series Services Gateways with a DOCSIS Mini-*Physical Interface Module* (Mini-PIM) located at customer premises and a cable modem termination system (CMTS) located at the head-end or data center locations. Standards-based DOCSIS 3.0 Mini-PIM is interoperable with CMTS equipment. The DOCSIS Mini-PIM provides backward compatibility with CMTS equipment based on the following standards:

- DOCSIS 2.0
- DOCSIS 1.1
- DOCSIS 1.0

The cable modem interface of Mini-PIM is managed and monitored by CMTS through SNMP. This DOCSIS 3.0 Mini-PIM can be deployed in any multiple service operator (MSO) networks. The primary application is for distributed enterprise offices to connect to a CMTS network through the DOCSIS 3.0 (backward compatible to 2.0, 1.1, and 1.0) interface. The DOCSIS Mini-PIM uses PIM infrastructure developed for third-party PIMs.

The Mini-PIM can also be used with encapsulations other than GRE, PPPoE, and IP-in-IP.

NOTE: The following interface trace options are supported:

- **all**—Enable all interface trace flags

- **event**—Trace interface events
- **ipc**—Trace interface IPC messages
- **media**—Trace interface media changes

CMTS manages and monitors the cable modem interface of then Mini-PIM through SNMP. This DOCSIS 3.0 Mini-PIM can be deployed in any multiple MSO network. Figure 29 on page 529 shows a typical use for this Mini-PIM in an MSO network.

Figure 29: Typical DOCSIS End-to-End Connectivity Diagram

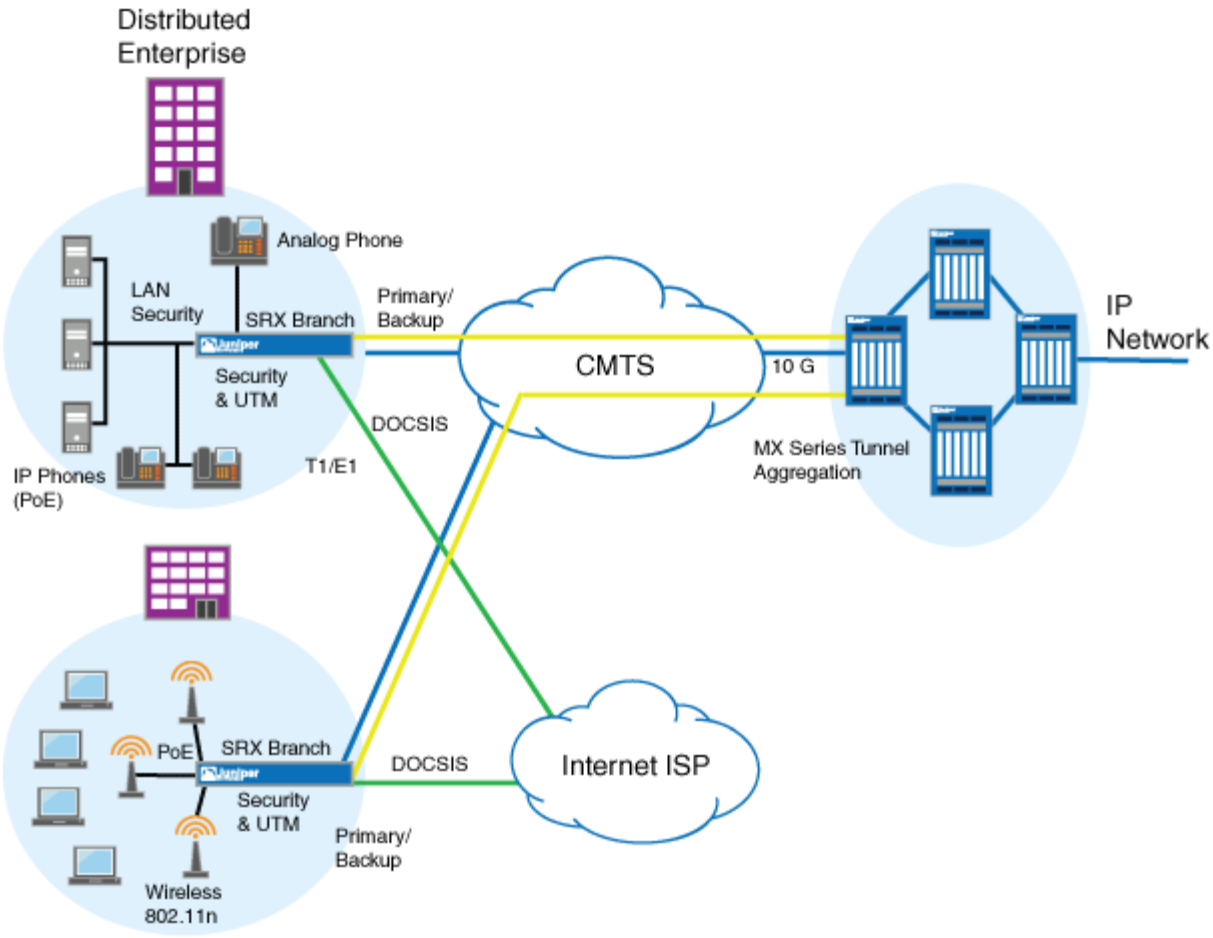

Software Features Supported on DOCSIS Mini-PIMs

NOTE: DOCSIS Mini-PIM interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Table 49 on page 530 lists the software features supported on DOCSIS Mini-PIMs.

Table 49: Software Features Supported on DOCSIS Mini-PIMs

Software Feature	Description
DHCP and DHCPv6 clients	The DHCP and DHCPv6 clients are used to get the IP address from the CMTS using the DHCP protocol. DHCP is supported on IPv4 and IPv6. One of the main components of the configuration file is the static public IP address, which CMTS assigns to the cable modem. The management IP address is configured on the Mini-PIM's hybrid fiber coaxial (HFC) interface, which performs the following tasks: • Allows CMTS to execute remote monitoring and management of the Mini-PIM's cable interface. • Downloads the configuration file from CMTS and uses it for configuring the cable interface.
QoS support	The SRX Series device's Routing Engine is configured through the existing QoS CLI. Because the configuration on the SRX Series device's Routing Engine and Mini-PIM is done together, the QoS configuration has to be consistent between the Routing Engine and the cable modem interface. The QoS mechanisms on the Routing Engine are decoupled from the QoS mechanisms on the Mini-PIM. The configuration file downloaded from CMTS contains parameters for primary and secondary flows. These parameters are programmed in the DOCSIS Mini-PIM. The Mini-PIM sends these parameters to the Routing Engine through the PIM infrastructure. The secondary flows are prioritized over primary flows in the DOCSIS Mini-PIM.

Table 49: Software Features Supported on DOCSIS Mini-PIMs *(Continued)*

Software Feature	Description
SNMP support	CMTS issues the SNMP requests that go to the cable modem. The DOCSIS MIB on the SRX Series device's Routing Engine displays the Ethernet interface of the cable modem. The following features are supported on the DOCSIS Mini-PIM: • NAT support • Dying gasp support • Back pressure information
MAC address	The MAC address of the DOCSIS Mini-PIM is statically set at the factory and cannot be changed. The MAC address is retrieved from the Mini-PIM and assigned to the cable modem interface in Junos OS.
Transparent bridging	The DOCSIS Mini-PIM performs transparent bridging by sending the packets received on the Ethernet interface with the SRX Series device to the HFC interface and vice versa, without any modifications to the packet. All the other services such as webserver, DHCP server, and DNS server are disabled on the DOCSIS Mini-PIM during transparent bridging.

Example: Configuring the DOCSIS Mini-PIM Interfaces

IN THIS SECTION

- Requirements | 532
- Overview | 532
- Configuration | 532
- Verification | 534

This example shows how to configure DOCSIS Mini-PIM network interfaces for SRX210, SRX220, and SRX240 devices.

NOTE: DOCSIS Mini-PIM interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM devices.

Requirements

Before you begin:

- Establish basic connectivity. See the Quick Start for your device.
- Configure network interfaces as necessary. See "Example: Creating an Ethernet Interface" on page 166.

Overview

In this example, you configure the DOCSIS Mini-PIM interface as cm-2/0/0. You specify the physical properties by setting the interface trace options and the flag option. You then set the logical interface to unit 0 and specify the family protocol type as inet. Finally, you configure the DHCP client.

Configuration

IN THIS SECTION

- Procedure | 532

Procedure

CLI Quick Configuration

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

```
set interfaces cm-2/0/0 traceoptions flag all
set interfaces cm-2/0/0 unit 0 family inet dhcp
```

Step-by-Step Procedure

To configure the DOCSIS Mini-PIM network interfaces:

1. Configure the interface.

```
[edit]
user@host# edit interfaces cm-2/0/0
```

2. Set the interface trace options.

```
[edit]
user@host# set interfaces cm-2/0/0 traceoptions
```

3. Specify the flag option.

```
[edit]
user@host# set interfaces cm-2/0/0 traceoptions flag all
```

4. Set the logical interface.

```
[edit]
user@host# set interfaces cm-2/0/0 unit 0
```

5. Specify the family protocol type.

```
[edit]
user@host# set interfaces cm-2/0/0 unit 0 family inet
```

6. Configure the DHCP client.

```
[edit]
user@host# set interfaces cm-2/0/0 unit 0 family inet dhcp
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces cm-2/0/0` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
    user@host# show interfaces cm-2/0/0
    traceoptions {
    flag all;
    }
    unit 0 {
        family inet {
    dhcp;
    }
}
```

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

Verification

IN THIS SECTION

- Verifying the DOCSIS Interface Properties | 534

Confirm that the configuration is working properly.

Verifying the DOCSIS Interface Properties

Purpose

Verify that the DOCSIS interface properties are configured properly.

Action

From operational mode, enter the `show interfaces cm-2/0/0` command.

```
user@host> show interfaces cm-2/0/0 extensive
Physical interface: cm-2/0/0, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 522, Generation: 157
  Link-level type: Ethernet, MTU: 1518, Speed: 40mbps
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
   State   : OPERATIONAL, Mode: 2.0, Upstream speed: 5120000 0 0 0
   Downstream scanning: CM_MEDIA_STATE_DONE, Ranging: CM_MEDIA_STATE_DONE
   Signal to noise ratio: 31.762909 21.390018  7.517472 14.924058
   Power: -15.756125 -31.840363 -31.840363 -31.840363
   Downstream buffers used           : 0
   Downstream buffers free           : 0
   Upstream buffers free             : 0
   Upstream buffers used             : 0
   Request opportunity burst         : 0 MSlots
  Physical burst                    : 0 MSlots
   Tuner frequency                   : 555 0 0 0 MHz
   Standard short grant              : 0 Slots
   Standard long grant               : 0 Slots
   Baseline privacy state: authorized, Encryption algorithm: ????, Key length: 0
  MAC statistics:                      Receive         Transmit
    Total octets                          1935             2036
    Total packets                            8                8
    CRC/Align errors                         0                0
    Oversized frames                         0
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:24:dc:0d:76:19, Hardware address: 00:24:dc:0d:76:19
  Last flapped   : 2009-11-10 19:55:40 UTC (00:16:29 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 710                    0 bps
   Output bytes  :                 866                    0 bps
   Input  packets:                   2                    0 pps
   Output packets:                   4                    0 pps
  Packet Forwarding Engine configuration:
    Destination slot: 1
    Direction : Output
    CoS transmit queue              Bandwidth              Buffer Priority   Limit
```

```
                            %           bps     %           usec
 0 best-effort             95      38000000    95              0      low    none
 3 network-control          5       2000000     5              0      low    none
Logical interface cm-2/0/0.0 (Index 69) (SNMP ifIndex 523) (Generation 134)
 Flags: Point-To-Point SNMP-Traps Encapsulation: ENET2
 Traffic statistics:
  Input  bytes  :                  710
  Output bytes  :                  806
  Input  packets:                    2
  Output packets:                    4
 Local statistics:
  Input  bytes  :                  710
  Output bytes  :                  806
  Input  packets:                    2
  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
 Security: Zone: Null
 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: 1504, Generation: 147, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 20.20.20/24, Local: 20.20.20.5, Broadcast: 20.20.20.255, Generation: 144
```

The output shows a summary of DOCSIS interface properties. Verify the following information:

- The physical interface is `Enabled`. If the interface is shown as `Disabled`, do either of the following:
  - In the CLI configuration editor, delete the `disable` statement at the `[edit interfaces` *`interface-name`*`]` level of the configuration hierarchy.
  - In the J-Web configuration editor, clear the `Disable` check box on the `Interfaces>`*`interface-name`* page.
- The physical link is `Up`. A link state of `Down` indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The `Last Flapped` time is an expected value. The `Last Flapped` time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect the expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches the expected throughput for the physical interface. To clear the statistics and see only new changes, use the `clear interfaces statistics` *`interface-name`* command.

Release History Table

Release	Description
15.1X49-D10	DOCSIS Mini-PIM interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

11

CHAPTER

Configuration Statements

accept-source-mac | 543

access-point name | 545

apply-groups | 546

activation-delay | 548

authentication-method (Interfaces) | 549

bandwidth (Interfaces) | 551

bundle (Interfaces) | 552

cbr rate | 553

callback | 555

callback-wait-period | 556

caller | 558

cellular-options | 560

classifiers (CoS) | 561

client-identifier (Interfaces) | 563

code-points (CoS) | 565

compression-device (Interfaces) | 567

credit (Interfaces) | 568

data-rate | 569

deactivation-delay | 571

disable (PoE) | 573

dialer-options | 574

dialin | 576

dial-string | 578

dhcp (DHCP Client) | 579

dsl-sfp-options | 582

duration (PoE) | 586

family inet (Interfaces) | 587

family inet6 | 592

flag (Interfaces) | 596

flexible-vlan-tagging (Interfaces) | 598

flow-control (Interfaces) | 599

flow-monitoring (Services) | 601

forwarding-classes | 603

fpc (Interfaces) | 609

gratuitous-arp-reply | 610

gsm-options | 612

guard-band (PoE) | 614

hold-time (Redundant Ethernet Interfaces) | 615

hub-assist | 617

idle-timeout | 619

incoming-map | 620

initial-route-check | 622

inline-jflow (Forwarding Options) | 623

interface (PIC Bundle) | 625

interface (PoE) | 627

interfaces (CoS) | 629

interval (Interfaces) | 631

interval (PoE) | 633

isdn-options | 634

ipv4-template (Services) | 636

ipv6-template (Services) | 637

lacp (Interfaces) | 639

latency (Interfaces) | 641

lease-time | 642

line-rate (Interfaces) | 644

link-speed (Interfaces) | 645

load-interval | 647

load-threshold | 648

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | 650

loss-priority (CoS Loss Priority) | 652

loss-priority (CoS Rewrite Rules) | 654

loss-priority-maps (CoS Interfaces) | 656

loss-priority-maps (CoS) | 657

management (PoE) | 659

maximum-power (PoE) | 661

mdi-mode | 662

media-type (Interfaces) | 665

minimum-links (Interfaces) | 667

modem-options | 668

mtu | 670

native-vlan-id | 671

next-hop-tunnel | 675

no-dns-propagation | 677

option-refresh-rate (Services) | 678

pic-mode (Chassis T1 Mode) | 680

periodic (Interfaces) | 681

pool | 683

ppp-over-ether | 685

pppoe | 686

pppoe-options | 688

priority (PoE) | 690

profile (Access) | 692

profiles | 697

promiscuous-mode (Interfaces) | 698

quality (Interfaces) | 700

r2cp | 701

radio-router (Interfaces) | 703

redial-delay | 705

redundancy-group (Interfaces) | 707

redundant-ether-options | 708

redundant-parent (Interfaces Fast Ethernet) | 711

redundant-parent (Interfaces Gigabit Ethernet) | 712

request pppoe connect | 714

request pppoe disconnect | 715

resource (Interfaces) | 717

retransmission-attempt (DHCP Client) | 719

retransmission-interval (DHCP Client) | 721

roaming-mode | 722

scheduler-map (CoS Virtual Channels) | 724

select-profile | 725

server-address | 727

shaping-rate (CoS Interfaces) | 729

simple-filter (Interfaces) | 731

sip-password | 732

sip-user-id | 734

source-address-filter (Interfaces) | 735

source-filtering (Interfaces) | 737

speed (Interfaces) | 739

speed (Gigabit Ethernet interface) | 740

spid1 | 742

spid2 | 743

static-tei-val | 745

switch-type | 746

t310 | 748

tei-option | 749

telemetries (PoE) | 751

template-refresh-rate (Services) | 753

threshold (Interfaces) | 754

traceoptions (Interfaces) | 756

update-server | 757

vbr rate | 759

vdsl-profile | 760

vendor-id (Interfaces) | 762

watch-list | 764

web-authentication (Interfaces) | 765

wlan | 767

accept-source-mac

IN THIS SECTION

- Syntax | 543
- Hierarchy Level | 543
- Description | 543
- Options | 544
- Required Privilege Level | 544
- Release Information | 544

Syntax

```
accept-source-mac {
    mac-address mac-address;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

For Gigabit Ethernet (GE), Fast Ethernet (FE), or 10 Gigabit Ethernet (XE) interfaces, specify the MAC addresses from which the interface can receive packets. Ensure that you update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. If you do not update the MAC address, the interface cannot receive packets from the new card.

NOTE:

- Software-based MAC limiting is supported on SRX300, SRX320, and SRX340 devices. A maximum of 32 MAC addresses is supported per device.

Options

mac-address —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* (for example, `00:11:22:33:44:55`) or *nnnn:nnnn:nnnn* (for example, `0011.2233.4455`). You can configure up to 32 source addresses. To specify more than one address, include multiple *mac-addresses* in the `source-address-filter` statement.

Required Privilege Level

interface—To view this statement in the configuration..

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

access-point name

IN THIS SECTION

- Syntax | 545
- Hierarchy Level | 545
- Description | 545
- Options | 546
- Required Privilege Level | 546
- Release Information | 546

Syntax

```
access-point-name apn;
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Configure the access point name (APN) provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

apn—Access point name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

apply-groups

IN THIS SECTION

- Syntax | 546
- Hierarchy Level | 547
- Description | 547
- Required Privilege Level | 547
- Release Information | 547

Syntax

```
apply-groups;
```

Hierarchy Level

```
[edit interfaces interface-name unit  logical-unit-number radio-router]
```

Description

Apply the groups from which to inherit configuration data. If `radio-router` is set without any other attributes specified, the first four values become 100 and threshold stays at 10, and capacity, margin, and delay are deprecated. If `radio-router` is set, do not change the OSPF reference-bandwidth value because this generates an incorrect link cost.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Statement modified in Junos OS Release 15.1.

RELATED DOCUMENTATION

Configuring PPPoE-Based Radio-to-Router Protocols

activation-delay

IN THIS SECTION

- Syntax | 548
- Hierarchy Level | 548
- Description | 548
- Options | 549
- Required Privilege Level | 549
- Release Information | 549

Syntax

```
activation-delay seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

(J Series Services Routers) For ISDN interfaces, configure the ISDN dialer activation delay. Used only for dialer backup and dialer watch cases.

Options

seconds—Interval before the backup interface is activated after the primary interface has gone down.

- **Range:** 1 through 4,294,967,295 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

authentication-method (Interfaces)

IN THIS SECTION

- Syntax | **550**
- Hierarchy Level | **550**
- Description | **550**
- Options | **550**
- Required Privilege Level | **550**
- Release Information | **550**

Syntax

```
authentication-method (pap | chap | none);
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Specify the authentication method for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

- pap—Password Authentication Protocol.
- chap—Challenge Handshake Authentication Protocol.
- none—No authentication method is used.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

bandwidth (Interfaces)

IN THIS SECTION

- Syntax | 551
- Hierarchy Level | 551
- Description | 551
- Required Privilege Level | 551
- Release Information | 552

Syntax

```
bandwidth bandwidth;
```

Hierarchy Level

```
[edit interfaces interface-name radio-router]
```

Description

This option controls the weight of the current (vs. maximum) data rate (value 0–100).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

- PPPoE-Based Radio-to-Router Protocols Overview

bundle (Interfaces)

IN THIS SECTION

- Syntax | 552
- Hierarchy Level | 552
- Description | 553
- Required Privilege Level | 553
- Release Information | 553

Syntax

```
bundle bundle-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family mlppp ]
```

Description

Specify the logical interface name the link joins.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces | 2

cbr rate

IN THIS SECTION

- Syntax | 554
- Hierarchy Level | 554
- Description | 554
- Options | 554
- Required Privilege Level | 554
- Release Information | 554

Syntax

```
cbr rate;
```

Hierarchy Level

```
[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
```

Description

For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.

Options

- CBR Value–Constant bandwidth utilization (range: 33,000 through 1,199,920)
- CDVT–Cell delay variation tolerance in microseconds (range: 1 through 9999)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Release 9.5 of Junos OS.

callback

IN THIS SECTION

- Syntax | 555
- Hierarchy Level | 555
- Description | 555
- Required Privilege Level | 556
- Release Information | 556

Syntax

```
callback;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options incoming-map],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options
incoming-map]
```

Description

On J Series Services Routers with interfaces configured for ISDN, configure the dialer to terminate the incoming call and call back the originator after the callback wait period. The default wait time is 5 seconds. To configure the wait time, include the `callback-wait-period` statement at the `[edit interfaces dl n unit logical-unit-number dialer-options]` hierarchy level.

NOTE: The `incoming-map` statement is mandatory for the router to accept any incoming ISDN calls.

If the `callback` statement is configured, you cannot use the `caller` *`caller-id`* statement at the `[edit interfaces dl`*`n`*` unit` *`logical-unit-number`* `dialer-options]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

callback-wait-period

callback-wait-period

IN THIS SECTION

- Syntax | 557
- Hierarchy Level | 557
- Description | 557
- Options | 557
- Required Privilege Level | 557
- Release Information | 558

Syntax

```
callback-wait-period time;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with interfaces configured for ISDN with callback, specify the amount of time the dialer waits before calling back the caller. The default wait time is 5 seconds. The wait time is necessary because, when a call is rejected, the switch waits for up to 4 seconds on point-to-multipoint connections to ensure no other device accepts the call before sending the DISCONNECT message to the originator of the call. However, the default time of 5 seconds may not be sufficient for different switches or may not be needed on point-to-point connections.

To configure callback mode, include the `callback` statement at the `[edit interfaces dln unit logical-unit-number dialer-options]` hierarchy level.

If the `callback` statement is configured, you cannot use the `caller caller-id` statement at the `[edit interfaces dln unit logical-unit-number dialer-options]` hierarchy level.

Options

time—Time the dialer waits before calling back the caller.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

caller

IN THIS SECTION

- Syntax | 558
- Hierarchy Level | 558
- Description | 559
- Options | 559
- Required Privilege Level | 559
- Release Information | 559

Syntax

```
caller (caller-id | accept-all);
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options incoming-map],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options
incoming-map]
```

Description

On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept a specified caller number or accept all incoming calls.

Options

caller-id—Incoming caller number. You can configure multiple caller IDs on a dialer. The caller ID of the incoming call is matched against all caller IDs configured on all dialers. The dialer matching the caller ID is looked at for further processing. Only a precise match is a valid match, For example, the configured caller ID 1-222-333-4444 or 222-333-4444 will match the incoming caller ID 1-222-333-4444.

If the incoming caller ID has fewer digits than the number configured, it is not a valid match. Duplicate caller IDs are not allowed on different dialers; however, for example, the numbers 1-408-532-1091, 408-532-1091, and 532-1091 can still be configured on different dialers.

Only one B-channel can map to one dialer. If one dialer is already mapped, any other call mapping to the same dialer is rejected (except in the case of a multilink dialer). If no dialer caller is configured on a dialer, that dialer will not accept any calls.

`accept-all`—Any incoming call in an associated interface is accepted.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

cellular-options

IN THIS SECTION

- Syntax | 560
- Hierarchy Level | 560
- Description | 561
- Options | 561
- Required Privilege Level | 561
- Release Information | 561

Syntax

```
cellular-options {
    roaming-mode (home only | automatic)
    gsm-options {
        select-profile profile-name;
        profiles {
            profile-name {
                sip-user-id simple-ip-user-id;
                sip-password simple-ip-password;
                access-point-name apn;
                authentication-method (pap | chap | none);
            }
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure options for connecting a 3G wireless modem interface to a cellular network.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

classifiers (CoS)

IN THIS SECTION

- Syntax | 562
- Hierarchy Level | 562
- Description | 562
- Options | 562
- Required Privilege Level | 563
- Release Information | 563

Syntax

```
classifiers {
    (dscp | dscp-ipv6 | exp | ieee-802.1 | ieee-802.1ad | inet-precedence) classifier-name {
        forwarding-class forwarding-class-name {
        loss-priority (high  | low  | medium-high  | medium-low) {
            code-point alias-or-bit-string ;
        }
        import (default | user-defined;
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Configure a user-defined behavior aggregate (BA) classifier.

Options

- *classifier-name*—User-defined name for the classifier.
- import (default | *user-defined*)—Specify the template to use to map any code points not explicitly mapped in this configuration. For example, if the classifier is of type `dscp` and you specify `import default`, code points you do not map in your configuration will use the predefined DSCP default mapping; if you specify `import mymap`, for example, code points not mapped in the forwarding-class configuration would use the mappings in a user-defined classifier named `mymap`.
- forwarding-class *class-name*—Specify the name of the forwarding class. You can use the default forwarding class names or define new ones.
- loss-priority *level*—Specify a loss priority for this forwarding class: `high`, `low`, `medium-high`, `medium-low`.

- code-points (*alias* | *bits*)—Specify a code-point alias or the code points that map to this forwarding class.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2

RELATED DOCUMENTATION

Understanding Interfaces

client-identifier (Interfaces)

IN THIS SECTION

- Syntax | 564
- Hierarchy Level | 564
- Description | 564
- Options | 564
- Required Privilege Level | 564
- Release Information | 565

Syntax

```
client-identifier {
    (ascii string | hexadecimal string);
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number
family family-name dhcp]
```

Description

Specify an ASCII or hexadecimal identifier for the Dynamic Host Configuration Protocol (DHCP) client. The DHCP server identifies a client by a client-identifier value.

Options

- `ascii` *ascii* —Identifier consisting of ASCII characters.
- `hexadecimal` *hexadecimal* —Identifier consisting of hexadecimal characters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces | 2

code-points (CoS)

IN THIS SECTION

- Syntax | 565
- Hierarchy Level | 565
- Description | 566
- Options | 566
- Required Privilege Level | 566
- Release Information | 566

Syntax

```
code-points ([ aliases ] | [ bit-patterns ]);
```

Hierarchy Level

```
[edit class-of-service classifiers type classifier-name forwarding-class class-name loss-
priority level]
```

Description

Specify one or more DSCP code-point aliases or bit sets to apply to a forwarding class..

NOTE: OCX Series switches do not support MPLS, and therefore, do not support EXP code points or code point aliases.

Options

aliases—Name of the DSCP alias.

bit-patterns—Value of the code-point bits, in six-bit binary form.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Understanding Interfaces

Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic

Example: Configuring Behavior Aggregate Classifiers

Example: Configuring BA Classifiers on Transparent Mode Security Devices

compression-device (Interfaces)

IN THIS SECTION

- Syntax | 567
- Hierarchy Level | 567
- Description | 567
- Options | 567
- Required Privilege Level | 568
- Release Information | 568

Syntax

```
compression-device name;
```

Hierarchy Level

```
[edit interfaces interface-name unit (Interfaces) logical-unit-number]
```

Description

Specify the compression interface for voice services traffic.

Options

name—Name of the AC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces | 2

credit (Interfaces)

IN THIS SECTION

- Syntax | 568
- Hierarchy Level | 569
- Description | 569
- Required Privilege Level | 569
- Release Information | 569

Syntax

```
credit {
    interval number;
}
```

Hierarchy Level

```
[edit interfaces interface-name radio-router ]
```

Description

This parameter controls credit-based scheduling parameters and includes an interval option to set the grant rate interval to a value between 1–60 seconds.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces | 2

data-rate

IN THIS SECTION

Syntax | 570

- Hierarchy Level | 570
- Description | 570
- Options | 570
- Required Privilege Level | 571
- Release Information | 571

Syntax

```
data-rate weight;
```

Hierarchy Level

```
[edit interfaces interface-name unit  logical-unit-number radio-router]
```

Description

Configure the weight of the resource factor when calculating an effective data rate.

Options

weight—Factor used to calculate data rate.

- **Range:** 0 through 100
- **Default:** 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.2 of Junos OS .

RELATED DOCUMENTATION

Configuring PPPoE-Based Radio-to-Router Protocols

deactivation-delay

IN THIS SECTION

- Syntax | 571
- Hierarchy Level | 572
- Description | 572
- Options | 572
- Required Privilege Level | 572
- Release Information | 572

Syntax

```
deactivation-delay seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, configure the ISDN deactivation delay. Used only for dialer backup and dialer watch cases.

Options

seconds—Interval before the backup interface is deactivated after the primary interface has comes up.

- **Range:** 1 through 4,294,967,295 seconds
- **Default:** 0 (zero)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

disable (PoE)

IN THIS SECTION

- Syntax | 573
- Hierarchy Level | 573
- Description | 573
- Default | 574
- Required Privilege Level | 574
- Release Information | 574

Syntax

```
disable;
```

Hierarchy Level

```
[edit poe interface  (all | interface-name) ]
 [edit poe interface (all | interface-name) telemetries]
```

Description

Disables the PoE capabilities of the port. If PoE capabilities are disabled for a port, the port operates as a standard network access port. If the disable statement is specified after the telemetries statement, logging of PoE power consumption for the port is disabled. To disable monitoring and retain the stored interval and duration values for possible future use, you can specify the disable sub statement in the sub stanza for telemetries. Similarly for retaining the port configuration but disabling the PoE feature on the port, disable can be used in sub stanza for interface.

Default

The PoE capabilities are automatically enabled when a PoE interface is set. Specifying the telemetries statement enables monitoring of PoE per-port power consumption.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

dialer-options

IN THIS SECTION

- Syntax | 574
- Hierarchy Level | 575
- Description | 575
- Required Privilege Level | 576
- Release Information | 576

Syntax

```
dialer-options {
    activation-delay seconds;
```

```
    callback;
    callback-wait-period time;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    incoming-map {
        caller caller-number | accept-all;
        initial-route-check seconds;
        load-interval seconds;
        load-threshold percent;
        pool pool-name;
        redial-delay time;
        self-recover-time
        watch-list {
            [ routes ];
        }
    }
}
```

Hierarchy Level

```
[edit interfaces umd0],
[edit interfaces dln unit logical-unit-number],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number]
```

Description

Specify the dialer options for configuring logical interfaces for group and user sessions.

The remaining statements are explained separately. See CLI Explorer.

You can use the new CLI option `self-recover-time` to configure the amount of time the `cl` interface waits to reconnect to a network after a disconnect occurs. In certain ISP networks, the modem disconnects and then reconnects after several seconds. The `self-recover-time` option provides enough time for the `cl` interface to reconnect instead of failing over to another `cl` interface immediately. If the reconnection attempt times out, then the connection fails over to another `cl` interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Junos OS Services Interfaces Library for Routing Devices

dialin

IN THIS SECTION

- Syntax | 576
- Hierarchy Level | 577
- Description | 577
- Options | 577
- Required Privilege Level | 577
- Release Information | 577

Syntax

```
dialin (console | routable);
```

Hierarchy Level

```
[edit interfaces umd0 modem-options]
```

Description

For J Series Services Routers, configure a USB modem port to act as a dial-in console or WAN backup port.

Options

console Configure the USB modem port to operate as a dial-in console for management.

routable Configure the USB modem port to operate as a dial-in WAN backup interface.

- **Default:** console

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

dial-string

IN THIS SECTION

- Syntax | 578
- Hierarchy Level | 578
- Description | 578
- Options | 579
- Required Privilege Level | 579
- Release Information | 579

Syntax

```
dial-string [ dial-string-numbers ];
```

Hierarchy Level

```
[edit interfaces br-pim/0/port unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces br-pim/0/port unit logical-unit-number
dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, specify one or more ISDN dial strings used to reach a destination subnetwork.

Options

dial-string-numbers—One or more strings of numbers to call.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

dhcp (DHCP Client)

IN THIS SECTION

- Syntax (EX Series) | **580**
- Syntax (SRX Series) | **580**
- Hierarchy level (EX Series) | **580**
- Hierarchy level (SRX Series) | **581**
- Description | **581**
- Options | **581**
- Required Privilege Level | **582**
- Release Information | **582**

Syntax (EX Series)

```
dhcp {
    client-identifier duid-type (duid-ll | duid-llt | vendor);
    no-dns-install;
    rapid-commit;
    options name;
}
```

Syntax (SRX Series)

```
dhcp {
    client-identifier {
        (ascii string | hexadecimal string);
    }
    force-discover;
    lease-time (length | infinite);
    metric;
    no-dns-install;
    options;
    requested-options;
    retransmission-attempt value;
    retransmission-interval seconds;
    server-address server-address;
    update-server;
    vendor-id vendor-id ;
}
```

Hierarchy level (EX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet]
[edit logical-systems name interfaces interface-name unit logical-unit-number family inet]
[edit tenants tenant-name interfaces interface-name unit logical-unit-number family inet]
```

Hierarchy level (SRX Series)

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Configure a Dynamic Host Configuration Protocol (DHCP) client for an IPv4 interface for logical systems and tenant systems.

The remaining statements are described separately.

NOTE: Starting in Junos OS Release 18.1R1, DHCPv4 and DHCPv6 clients are supported on management interfaces (fxp0 and em0) configured in the non-default management routing instance, `mgmt_junos`.

Options

client-identifier duid-type	Identify a client by a client-identifier value. This statement is mandatory.
no-dns-install	Do not add DNS information to the DHCP client even after it is learned from the DHCP server.
options	Specify options requested by the DHCPv4 client.
force-discover	Send DHCPDISCOVER after DHCPREQUEST retransmission failure
lease-time	Specify lease time in seconds requested in DHCP client protocol packet (60 through 2147,,483,647 seconds for SRX devices)
metric	client initiated default-route metric (0..255 for SRX Series devices)
requested-options	Specify the DHCP options.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

The `logical-systems` and `tenants` options are introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

- Understanding Interfaces
- *Configuring a DHCP Client*

dsl-sfp-options

IN THIS SECTION

- Syntax (SRX300, SRX320, SRX340, SRX345, SRX380) | **583**
- Hierarchy Level | **583**
- Description | **583**
- Options | **584**
- Required Privilege Level | **585**
- Release Information | **585**

Syntax (SRX300, SRX320, SRX340, SRX345, SRX380)

```
dsl-sfp-options {
    adsl-options {
        encap encapsulation;
        vci identifier;
        vpi identifier;
        annex  (auto | annexj-off);
    }
    gfast-options{
        carrier carrier setting;
    }
    vdsl-options {
        carrier carrier setting;
        profile profile;
    }
}
```

Hierarchy Level

```
[edit interfaces interface name]
```

Description

Configure ADSL properties on SRX Series devices.

Follow these example pre-configuration steps with the SFP inserted on slot 8 of the device, required for Annex J support. You must configure one logical interface for *x*DSL SFP control path to work.

1. ADSL and Annex J on `inet` interface:

```
user@host# set interfaces ge-0/0/8 unit 0 family inet address 10.1.1.10/24
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 encap llcsnap-
bridged-802.1q
```

```
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 annex auto encap
llcsnap-bridged-802.1q
```

2. ADSL and Annex J without vlan on Ethernet-switching interface:

```
user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching vlan members xdsl-test
user@host# set interfaces irb unit 50 family inet address 10.1.1.10/24
user@host# set vlans xdsl-test vlan-id 50
user@host# set vlans xdsl-test l3-interface irb.50
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 encap llcsnap-
bridged-802.1q
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 annex auto encap
llcsnap-bridged-802.1q
```

3. ADSL and Annex J with vlan on Ethernet-switching interface:

```
user@host# set interfaces ge-0/0/8 native-vlan-id 50
user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching interface-mode trunk
user@host# set interfaces ge-0/0/8 unit 0 family ethernet-switching vlan members xdsl-test
user@host# set interfaces irb unit 50 family inet address 10.1.1.10/24
user@host# set vlans xdsl-test vlan-id 50
user@host# set vlans xdsl-test l3-interface irb.50
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 encap llcsnap-
bridged-802.1q
user@host# set interfaces ge-0/0/8 dsl-sfp-options adsl-options vpi 8 vci 36 annex auto encap
llcsnap-bridged-802.1q
```

Similarly, to set up G.fast option under `dsl-sfp-options`, use:

```
user@host# set interfaces ge-0/0/8 unit 0 family inet address 10.1.1.1/24
user@host# set interfaces ge-0/0/8 dsl-sfp-options gfast-options carrier a43
```

Options

adsl-options Options of ADSL interface.

annex (auto | annexj-off)	The Annex type.
carrier	Carrier setting on VDSL. Supported VDSL carriers are A43, B43, and auto. Supported G.fast carriers are A43, A43c, B43, and B43c.
encap	The encapsulation type.
gfast-options	Options of G.fast (Gigabit broadband access technology).
profile	The profile type on VDSL.
vci	Virtual circuit identifier.
vdsl-options	Options.
vpi	Virtual path identifier.

Required Privilege Level

routing—To view this statement in the configuration.

routing-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.

Release Information

Statement introduced in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

Configuring ADSL Interfaces

duration (PoE)

IN THIS SECTION

- Syntax | 586
- Hierarchy Level | 586
- Description | 586
- Options | 587
- Required Privilege Level | 587
- Release Information | 587

Syntax

```
duration hours;
```

Hierarchy Level

```
[edit poe interface (all | interface-name) telemetries]
```

Description

Modifies the duration for which telemetry records are stored. If telemetry logging continues beyond the specified duration, the older records are discarded one by one as new records are collected.

Options

hours— Hours for which telemetry data should be retained.

- **Range:** 1 through 24 hours
- **Default:** 1 hour

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

family inet (Interfaces)

IN THIS SECTION

- Syntax | **588**
- Hierarchy Level | **590**
- Description | **591**
- Options | **591**
- Required Privilege Level | **591**
- Release Information | **591**

Syntax

```
inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
    address (source-address/prefix) {
        arp destination-address {
            (mac mac-address | multicast-mac multicast-mac-address);
            publish publish-address;
        }
        broadcast address;
        preferred;
        primary;
        vrrp-group group-id {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            advertisements-threshold number;
            authentication-key key-value;
            authentication-type (md5 | simple);
            fast-interval milliseconds;
            inet6-advertise-interval milliseconds
            (preempt <hold-time seconds> | no-preempt );
            priority value;
            track {
                interface interface-name {
                    bandwidth-threshold bandwidth;
                     priority-cost value;
                }
                priority-hold-time seconds;
                route route-address{
                    routing-instance routing-instance;
                     priority-cost value;
                }
            }
            virtual-address [address];
            virtual-link-local-address address;
```

```
            vrrp-inherit-from {
                active-group value;
                active-interface interface-name;
            }
        }
        web-authentication {
            http;
            https;
            redirect-to-https;
        }
    }
    dhcp {
        client-identifier {
            (ascii string | hexadecimal string);
        }
        lease-time (length | infinite);
        retransmission-attempt value;
        retransmission-interval seconds;
        server-address server-address;
        update-server;
        vendor-id vendor-id ;
    }
    dhcp-client {
        client-identifier {
            prefix {
                host-name;
                logical-system-name;
                routing-instance-name;
            }
            use-interface-description (device | logical);
            user-id (ascii string| hexadecimal string);
        }
        lease-time (length | infinite);
        retransmission-attempt value;
        retransmission-interval seconds;
        server-address server-address;
        update-server;
        vendor-id vendor-id ;
    }
    filter {
        group number;
        input filter-name;
        input-list [filter-name];
```

```
        output filter-name;
        output-list [filter-name];
    }
    mtu value;
    no-neighbor-learn;
    no-redirects;
    policer {
        arp arp-name;
        input input-name;
        output output-name;
    }
    primary;
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
    sampling {
        input;
        output;
        simple-filter;
    }
    targeted-broadcast {
        (forward-and-send-to-re |forward-only);
    }
    unnumbered-address {
        interface-name;
        preferred-source-address preferred-source-address;
    }
}
```

Hierarchy Level

```
[edit interfaces interface unit unit ]
```

Description

Assign an IP address to a logical interface.

Options

ipaddress—Specify the IP address for the interface. The remaining statements are explained separately.

> **NOTE**: You use family `inet` to assign an IPv4 address. You use family `inet6` to assign an IPv6 address. An interface can be configured with both an IPv4 and IPv6 address.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement supported in Junos 10.2 for SRX Series devices.

RELATED DOCUMENTATION

Understanding Interfaces | 2

family inet6

IN THIS SECTION

- Syntax | 592
- Hierarchy Level | 594
- Description | 595
- Options | 595
- Required Privilege Level | 595
- Release Information | 595

Syntax

```
inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            ouput;
        }
    }
    address source-address/prefix {
        eui-64;
        ndp address {
            (mac mac-address | multicast-mac multicast-mac-address);
            publish;
        }
        preferred;
        primary;
        vrrp-inet6-group group_id {
            (accept-data | no-accept-data);
            advertisements-threshold number;
            authentication-key value;
            authentication-type (md5 | simple);
            fast-interval milliseconds;
```

```
            inet6-advertise-interval milliseconds;
            (preempt <hold-time seconds>| no-preempt );
            priority value;
            track {
                interface interface-name {
                     bandwidth-threshold value;
                    priority-cost value;
                }
                priority-hold-time seconds;
                route route-address{
                    routing-instance routing-instance;
                }
            }
            virtual-inet6-address [address];
            virtual-link-local-address address;
            vrrp-inherit-from {
                active-group value;
                active-interface interface-name;
            }
        }
        web-authentication {
            http;
            https;
            redirect-to-https;
        }
    }
    (dad-disable | no-dad-disable);
    dhcpv6-client {
        client-ia-type (ia-na | ia-pd);
        client-identifier duid-type (duid-ll | duid-llt | vendor);
        client-type (autoconfig | stateful);
        rapid-commit;
        req-option (dns-server | domain | fqdn | nis-domain | nis-server | ntp-server | sip-
domain | sip-server |time-zone | vendor-spec);
        retransmission-attempt number;
        update-router-advertisement {
            interface interface-name;
        }
        update-server;
    }
    filter {
        group number;
        input filter-name;
```

```
        input-list [filter-name];
        output filter-name;
        output-list [filter-name];
    }
    mtu value;
    nd6-stale-time seconds;
    no-neighbor-learn;
    policer {
        input input-name;
        output output-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
    sampling {
        input;
        output;
    }
    unnumbered-address {
        interface-name;
        preferred-source-address preferred-source-address;
    }
    ndp-proxy | dad-proxy {
        interface-restricted
        }
}
```

Hierarchy Level

```
[edit interfaces interface unit unit ]
```

Description

Assign an IPV6 address to a logical interface.

Options

ipaddress—Specify the IP address for the interface. The remaining statements are explained separately.

> **NOTE**: You use family `inet6` to assign an IPv6 address. You use family `inet` to assign an IPv4 address. An interface can be configured with both an IPv4 and IPv6 address.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement supported in Junos 10.2 for SRX Series devices.

RELATED DOCUMENTATION

Understanding Interfaces | 2

flag (Interfaces)

IN THIS SECTION

- Syntax | 596
- Hierarchy Level | 596
- Description | 596
- Options | 597
- Required Privilege Level | 597
- Release Information | 597

Syntax

```
flag
```

Hierarchy Level

```
[edit interfaces interface-name traceoptions]
```

Description

Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.

Options

- **all**—Enable all interface trace flags.
- **event** —Trace interface events.
- **cache**—Enable interface flags for Web filtering cache maintained on the routing table.
- **enhanced**—Enable interface flags for processing through Enhanced Web Filtering.
- **ipc**—Trace interface IPC messages.
- **media**—Trace interface media changes.
- **critical**—Trace critical events.
- **major**—Trace major events.

NOTE:

- MTU is limited to 1518 on this interface.
- **Cache** and **enhanced** options are applicable only to Enhanced Web Filtering.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces | 2

flexible-vlan-tagging (Interfaces)

IN THIS SECTION

- Syntax | 598
- Hierarchy Level | 598
- Description | 598
- Options | 599
- Required Privilege Level | 599
- Release Information | 599

Syntax

```
flexible-vlan-tagging;
```

Hierarchy Level

```
[edit interfaces interface ]
```

Description

Simultaneously supports transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

NOTE: The `flexible-vlan-tagging` is supported only with either no encapsulation or VPLS VLAN encapsulation.

Options

native-vlan-id—Configures a VLAN identifier for single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

Configuring VLAN Tagging | 41

flow-control (Interfaces)

IN THIS SECTION

- Syntax | 600
- Hierarchy Level | 600
- Description | 600
- Default | 600
- Required Privilege Level | 600
- Release Information | 600

Syntax

```
(flow-control | no-flow-control);
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options]
[edit interfaces interface-name gigether-options]
[edit interfaces interface-name redundant-ether-options]
```

Description

For Fast Ethernet, Gigabit Ethernet, and redundant Ethernet interfaces, flow control regulates the flow of packets from the device to the remote side of the connection.

Default

Flow control is the default behavior for Fast Ethernet and Gigabit Ethernet interfaces. Flow control is disabled by default for redundant Ethernet interfaces

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

flow-monitoring (Services)

IN THIS SECTION

- Syntax | 601
- Hierarchy Level | 602
- Description | 602
- Options | 602
- Required Privilege Level | 602
- Release Information | 602

Syntax

```
flow-monitoring {
    version9 {
        template template-name {
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            ipv4-template;
            ipv6-template;
            option-refresh-rate {
                packets packets;
                seconds seconds;
            }
            template-refresh-rate {
                packets packets;
                seconds seconds;
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure flow monitoring.

Options

`version9`—Version 9 configuration.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Interfaces | 2

forwarding-classes

IN THIS SECTION

- SRX Series | 603
- QFX Series and OCX Series | 604
- EX Series (Except EX4300) | 604
- EX4300 | 604
- M320, MX Series, T Series, and PTX Series | 604
- Hierarchy Level | 605
- Description | 605
- Options | 607
- Required Privilege Level | 608
- Release Information | 608

SRX Series

```
forwarding-classes {
    class class-name {
        priority (high | low);
        queue-num number;
        spu-priority (high | low | medium);
    }
    queue queue-number {
        class class-name {
            priority (high | low);
        }
    }
}
```

QFX Series and OCX Series

```
forwarding-classes {
    class class-name {
        pfc-priority pfc-priority;
        no-loss;
        queue-num queue-number <no-loss>;
    }
}
```

EX Series (Except EX4300)

```
forwarding-classes {
    class class-name {
        queue-num queue-number;
        priority (high | low);
    }
}
```

EX4300

```
forwarding-classes {
    class class-name ;
        queue-num queue-number;
    }
}
```

M320, MX Series, T Series, and PTX Series

```
forwarding-classes {
    class class-name {
```

```
        queue queue-number;
        priority (high | low);
    }
    queue queue-number {
        class class-name {
            priority (high | low) [policing-priority (premium | normal)];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Command used to associate forwarding classes with class names and queues with queue numbers.

SRX Series Devices

All traffic traversing the SRX Series device is passed to an SPC to have service processing applied. Junos OS provides a configuration option to enable packets with specific Differentiated Services (DiffServ) code points (DSCP) precedence bits to enter a high-priority queue, a medium-priority queue, or a low-priority queue on the SPC. The Services Processing Unit (SPU) draws packets from the highest priority queue first, then from the medium priority queue, and last from the low priority queue. The processing of the queue is weighted-based not strict-priority-based. This feature can reduce overall latency for real-time traffic, such as voice traffic.

Initially, the spu-priority queue options were "high" and "low". Then, these options (depending on the devices) were expanded to "high", "medium-high", "medium-low", and "low". The two middle options ("medium-high" and "medium-low") have now been deprecated (again, depending on the devices) and replaced with "medium". So, the available options for spu-priority queue are "high", "medium", and "low".

We recommend that the high-priority queue be selected for real-time and high-value traffic. The other options would be selected based on user judgement on the value or sensitivity of the traffic.

M320, MX Series, and T Series Routers and EX Series Switches

For M320, MX Series, and T Series routers, and EX Series switches only, you can configure fabric priority queuing by including the `priority` statement. For Enhanced IQ PICs, you can include the `policing-priority` option.

NOTE: The `priority` and `policing-priority` options are not supported on PTX Series routers.

EX Series Switches

For the EX Series switches, this statement associates the forwarding class with a class name and queue number. It can define the fabric queuing priority as high, medium-high, medium-low, or low.

Map one or more forwarding classes to a single output queue. Also, when configuring DSCP-based priority-based flow control (PFC), map a forwarding class to a PFC priority value to use in pause frames when traffic on a DSCP value becomes congested (see *Configuring DSCP-based PFC for Layer 3 Untagged Traffic* for details).

Switches that use different forwarding classes for unicast and multidestination (multicast, broadcast, and destination lookup fail) traffic support 12 forwarding classes and 12 output queues (0 through 11). You map unicast forwarding classes to a unicast queue (0 through 7) and multidestination forwarding classes to a multidestination queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multidestination forwarding class.

Switches that use the same forwarding classes for unicast and multidestination traffic support eight forwarding classes and eight output queues (0 through 7). You map forwarding classes to output queues. All traffic classified into one forwarding class (unicast and multidestination) uses the same output queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the `no-loss` packet drop attribute. Do not associate a drop profile with lossless forwarding classes.

NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

OCX Series Switches

OCX Series switches do not support the `no-loss` packet drop attribute and do not support lossless forwarding classes. On OCX Series switches, do not configure the `no-loss` packet drop attribute on forwarding classes, and do not map traffic to the default `fcoe` and `no-loss` forwarding classes (both of these default forwarding classes carry the `no-loss` packet drop attribute).

NOTE: On switches that do not use the Enhanced Layer 2 Software (ELS) CLI, if you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless `fcoe` and `no-loss` forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

NOTE: On switches that do not use the ELS CLI, if you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the `fcoe` and `no-loss` forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the `no-loss` option. If you do not specify the `no-loss` option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

Options

class *class-name* — Define the forwarding class name.

queue-num *queue-number* — Output queue number to associate with forwarding class.

- **Range:** 0 through 7.

priority — Fabric priority value:

- **high** — Forwarding class fabric queuing has high priority.
- **low** — Forwarding class fabric queuing has low priority.

  The `default` priority is `low`.

spu-priority — SPU priority queue, `high`, `medium`, or `low`. The default `spu-priority` is `low`.

NOTE: The `spu-priority` option is supported only on the SRX5000 line of devices.

The remaining statements are explained separately. See CLI Explorer for details.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

The `policing-priority` option was introduced in Junos OS Release 9.5.

Statement updated in Junos OS Release 11.4.

The `spu-priority` option was introduced in Junos OS Release 11.4R2.

The `no-loss` option was introduced in Junos OS Release 12.3 on QFX Series switches.

Change from two to four queues made in Junos OS Release 12.3X48-D40 and in Junos OS Release 15.1X49-D70.

The `pfc-priority` statement was introduced in Junos OS Release 17.4R1 on QFX Series switches.

The `medium-high` and `medium-low` priorities for `spu-priority` were deprecated and `medium` priority was added in Junos OS Release 19.1R1.

RELATED DOCUMENTATION

- *Configuring a Custom Forwarding Class for Each Queue*
- *Forwarding Classes and Fabric Priority Queues*
- *Configuring Hierarchical Layer 2 Policers on IQE PICs*
- *Classifying Packets by Egress Interface*

fpc (Interfaces)

IN THIS SECTION

- Syntax | 609
- Hierarchy Level | 609
- Description | 609
- Options | 610
- Required Privilege Level | 610
- Release Information | 610

Syntax

```
fpc slot-number ;
```

Hierarchy Level

```
[edit interfaces pic-set pic-set-name]
```

Description

Sets the PIC bundle and the FPC slot.

The `pic-set` bundles all the PICs and corresponding logical interfaces. A PIC can only join only one pic-bundle, and cannot join multiple pic-bundles at same time. When the `pic-set` configuration changes, all the logical interfaces related to the PIC should be synchronized to all member IOC.

Options

- *apply-groups*—Inherit configuration data from these groups.
- *apply-groups-except*—Do not inherit configuration data from these groups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Interfaces | 2

interface (PIC Bundle) | 625

gratuitous-arp-reply

IN THIS SECTION

- Syntax | 611
- Hierarchy Level | 611
- Description | 611
- Default | 611
- Required Privilege Level | 611
- Release Information | 611

Syntax

```
(gratuitous-arp-reply | no-gratuitous-arp-reply);
```

Hierarchy Level

```
[edit interfaces interface-name]
[edit interfaces interface-range interface-range-name]
```

Description

For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.

Default

Updating of the ARP cache is disabled on all Ethernet interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

Configuring Gratuitous ARP

no-gratuitous-arp-request

gsm-options

IN THIS SECTION

- Syntax | 612
- Hierarchy Level | 613
- Description | 613
- Options | 613
- Required Privilege Level | 613
- Release Information | 613

Syntax

```
gsm-options {
    select-profile profile-name;
    profiles {
        profile-name {
            sip-user-id simple-ip-user-id;
            sip-password simple-ip-password;
            access-point-name apn;
            authentication-method (pap | chap | none);
        }
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options]
```

Description

Configure the 3G wireless modem interface to establish a data call with a Global System for Mobile Communications (GSM) cellular network.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

guard-band (PoE)

IN THIS SECTION

- Syntax | 614
- Hierarchy Level | 614
- Description | 614
- Options | 615
- Required Privilege Level | 615
- Release Information | 615

Syntax

```
guard-band watts;
```

Hierarchy Level

```
[edit poe]
```

Description

Reserves the specified amount of power for the SRX Series device in case of a spike in PoE consumption.

Options

watts—Amount of power to be reserved for the SRX Series device in case of a spike in PoE consumption.

- **Range:** 0 through 19 W
- **Default:** 0 W

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Power over Ethernet Overview | 261

hold-time (Redundant Ethernet Interfaces)

IN THIS SECTION

- Syntax | 616
- Hierarchy Level | 616
- Description | 616
- Options | 617
- Required Privilege Level | 617

- Release Information | 617

Syntax

```
hold-time (up | down) timer
```

Hierarchy Level

```
[edit interfaces interface-name ]
```

Description

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 for a parent RETH interface and the primary child interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the down hold-time is ignored. When the timer expires and the primary child interface state is still down, then the router begins to advertise the parent RETH interface as being down. Similarly, when a hold-up timer is configured for a parent RETH interface and the primary child interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the up hold-time is ignored. When the timer expires and the primary child interface state is still up, then the router begins to advertise the parent RETH interface as being up.

The hold timer (both up and down) improves the flexibility and resilience of SRX devices. Specify the *timer* value in seconds to reduce unnecessary loss of traffic and downtime.

NOTE: Starting in Junos OS release 18.4R1, all SRX devices have default delay timer of 11 seconds for both up hold-time and down hold-time.

Options

`down` *seconds*—Hold time to use when an interface transitions from up to down.

`up` *seconds*—Hold time to use when an interface transitions from down to up.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

Physical Interface Damping Overview

hold-time

hub-assist

IN THIS SECTION

- Syntax | **618**
- Hierarchy Level | **618**
- Description | **618**
- Options | **618**
- Required Privilege Level | **618**
- Release Information | **618**

Syntax

```
hub-assist weight;
```

Hierarchy Level

```
[edit interfaces interface-name unit  logical-unit-number radio-router]
```

Description

Configure the weight of the resource factor when calculating an effective interface bandwidth.

Options

weight—Factor used to calculate interface bandwidth.

- **Range:** 0 through 100
- **Default:** 100

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Configuring PPPoE-Based Radio-to-Router Protocols

idle-timeout

IN THIS SECTION

- Syntax | 619
- Hierarchy Level | 619
- Description | 619
- Options | 620
- Required Privilege Level | 620
- Release Information | 620

Syntax

```
idle-timeout seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, configure the number of seconds the link is idle before losing connectivity.

Options

seconds—Time for which the connection can remain idle. For interfaces configured to use a filter for traffic, the idle timeout is based on traffic.

- **Range:** 1 through 429497295
- **Default:** 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

incoming-map

IN THIS SECTION

- Syntax | **621**
- Hierarchy Level | **621**
- Description | **621**
- Required Privilege Level | **621**
- Release Information | **621**

Syntax

```
incoming-map {
    caller caller-number | accept-all;
}
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept incoming calls.

The remaining statements are explained separately. See CLI Explorer.

NOTE: The `incoming-map` statement is mandatory for the router to accept any incoming ISDN calls.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

initial-route-check

IN THIS SECTION

- Syntax | 622
- Hierarchy Level | 622
- Description | 622
- Options | 623
- Required Privilege Level | 623
- Release Information | 623

Syntax

```
initial-route-check seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, allows the router to check whether the primary route is up after the initial startup of the router is complete and the timer expires.

Options

seconds—How long to wait to check if the primary interface is up after the router comes up.

- **Range:** 1 through 300 seconds
- **Default:** 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

ISDN Interfaces Overview

inline-jflow (Forwarding Options)

IN THIS SECTION

- Syntax | **624**
- Hierarchy Level | **624**
- Description | **624**
- Options | **624**
- Required Privilege Level | **624**

Release Information | 625

Syntax

```
inline-jflow {
    flow-export-rate number;
    source-address ip-address;
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance- name family inet output]
[edit forwarding-options sampling instance instance- name family inet6 output]
```

Description

Specify Inline processing of sampled packets.

Options

- `flow-export-rate` *value*—Flow export rate of monitored packets in kpps. The range is from 1 through 400.
- `source-address` *address*—Address to use for generating monitored packets.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4. Support.

RELATED DOCUMENTATION

Understanding Interfaces | 2

interface (PIC Bundle)

IN THIS SECTION

- Syntax | 625
- Hierarchy Level | 626
- Description | 626
- Options | 626
- Required Privilege Level | 626
- Release Information | 626

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit interfaces pic-set pic-set-name]
```

Description

Sets the PIC bundle and the interface.

Options

- *apply-groups*– Groups from which to inherit configuration data.
- *apply-groups-except*– Do not inherit configuration data from these groups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

Understanding Interfaces | 2

interface (PoE)

IN THIS SECTION

- Syntax | 627
- Hierarchy Level | 627
- Description | 628
- Default | 628
- Options | 628
- Required Privilege Level | 628
- Release Information | 628

Syntax

```
interface  (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
        disable;
        duration hours;
        interval minutes;
    }
}
```

Hierarchy Level

```
[edit poe]
```

Description

Enable a PoE interface for a PoE port. The PoE interface must be enabled in order for the port to provide power to a connected powered device.

Default

The PoE interface is enabled by default

Options

- `all`— Apply the configuration to all interfaces on the SRX Series device that have not been explicitly configured otherwise.
- `interface-name`— Explicitly configure a specific interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Power over Ethernet Overview | 261

interfaces (CoS)

IN THIS SECTION

- Syntax | 629
- Hierarchy Level | 630
- Description | 630
- Options | 630
- Required Privilege Level | 630
- Release Information | 631

Syntax

```
interfaces
                ***interface-name*** {
                    input-scheduler-map map-name;
                    input-shaping-rate rate;
                    scheduler-map map-name;
                    scheduler-map-chassis map-name;
                    shaping-rate rate;
                    unit  logical-unit-number  {
                        adaptive-shaper adaptive-shaper-name;
                        classifiers {
                            dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence)
                            (classifier-name | default);
                        }
                        forwarding-class class-name;
                        fragmentation-map map-name;
                        input-scheduler-map map-name;
                        input-shaping-rate (percent percentage  |  rate);
                        input-traffic-control-profile profiler-name  shared-instance instance-
name;
                        loss-priority-maps {
                            default;
                            map-name;
```

```
                        }
                        output-traffic-control-profile profile-name shared-instance instance-
name;
                        rewrite-rules {
                            dscp (rewrite-name | default);
                            dscp-ipv6 (rewrite-name | default);
                            exp (rewrite-name | default) protocol protocol-types;
                            frame-relay-de (rewrite-name | default);
                            inet-precedence (rewrite-name | default);
                        }
                        scheduler-map map-name;
                        shaping-rate rate;
                        virtual-channel-group group-name;
                    }
                }
            }
```

Hierarchy Level

```
[edit class-of-service interface interface-name unit number]
```

Description

Associate the class-of-service configuration elements with an interface.

Options

interface *interface-name* unit *number*—The user-specified interface name and unit number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

Class of Service User Guide (Security Devices)

interval (Interfaces)

IN THIS SECTION

- Syntax | 631
- Hierarchy Level | 632
- Description | 632
- Options | 632
- Required Privilege Level | 632
- Release Information | 632

Syntax

```
interval seconds;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number radio-router credit]
```

Description

Configure the frequency that the router generates credit announcement messages.

Options

seconds—Interval between PADG credit announcements for each session.

- **Range:** 0 through 60
- **Default:** 1

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 10.1 of Junos OS.

RELATED DOCUMENTATION

Configuring PPPoE-Based Radio-to-Router Protocols

interval (PoE)

IN THIS SECTION

- Syntax | 633
- Hierarchy Level | 633
- Description | 633
- Options | 634
- Required Privilege Level | 634
- Release Information | 634

Syntax

```
interval minutes;
```

Hierarchy Level

```
[edit poe interface (all | interface-name) telemetries]
```

Description

Modifies the interval for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.

Options

minutes—Interval at which data is logged.

- **Range:** 1 through 30 minutes
- **Default:** 5 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces | 2

isdn-options

IN THIS SECTION

- Syntax | 635
- Hierarchy Level | 635
- Description | 635
- Required Privilege Level | 635
- Release Information | 636

Syntax

```
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    incoming-called-number number <reject>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
```

Hierarchy Level

```
[edit interfaces br-pim/0/port],
[edit interfaces ct1-pim/0/port],
[edit interfaces ce1-pim/0/port]
```

Description

For J Series Services Routers only. Specify the ISDN options for configuring ISDN interfaces for group and user sessions.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`bchannel-allocation` option added in Junos OS Release 8.3.

RELATED DOCUMENTATION

- Configuring ISDN Physical Interface Properties
- Allocating B-Channels for Dialout

ipv4-template (Services)

IN THIS SECTION

- Syntax | 636
- Hierarchy Level | 636
- Description | 637
- Required Privilege Level | 637
- Release Information | 637

Syntax

```
ipv4-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specify that the flow monitoring version 9 template is used only for IPv4 records.

Required Privilege Level

services—To view this in the configuration.

services-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

Understanding Interfaces | 2

ipv6-template (Services)

IN THIS SECTION

- Syntax | 638
- Hierarchy Level | 638
- Description | 638
- Required Privilege Level | 638
- Release Information | 638

Syntax

```
ipv6-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specify that the flow monitoring version 9 template is used only for IPv6 records.

Required Privilege Level

services—To view this in the configuration.

services-control—To add this to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

Understanding Interfaces | 2

lacp (Interfaces)

IN THIS SECTION

- Syntax | 639
- Hierarchy Level | 639
- Description | 639
- Options | 640
- Required Privilege Level | 640
- Release Information | 640

Syntax

```
lacp {
    (active | passive);
    periodic (fast | slow);
}
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces in a chassis cluster only, configure Link Aggregation Control Protocol (LACP).

Options

- `active`—Initiate transmission of LACP packets.
- `passive`—Respond to LACP packets.
- `periodic`—Interval for periodic transmission of LACP packets. The options are:
  - `fast`—Transmit link aggregation control PDUs every second.
  - `slow`—Transmit link aggregation control PDUs every 30 seconds.
- **Default:** If you do not specify `lacp` as either `active` or `passive`, LACP remains off (the default).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding LACP on Standalone Devices

periodic (Interfaces)

latency (Interfaces)

IN THIS SECTION

- Syntax | 641
- Hierarchy Level | 641
- Description | 641
- Required Privilege Level | 641
- Release Information | 642

Syntax

```
latency number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number radio-router ]
```

Description

This option controls the latency weight (value 0–100).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

lease-time

IN THIS SECTION

- Syntax | 642
- Hierarchy Level | 643
- Description | 643
- Default | 643
- Options | 643
- Required Privilege Level | 643
- Release Information | 644

Syntax

```
lease-time (length | infinite);
```

Hierarchy Level

```
[edit interfaces interface-name                unit                logical-unit-
number            family     inet dhcp]
```

Description

Request a specific lease time for the IP address. The lease time is the length of time in seconds that a client holds the lease for an IP address assigned by a DHCP server.

Default

If no lease time is requested by client, then the server sends the lease time. The default lease time on a Junos OS DHCP server is one day.

Options

seconds Request a lease time of a specific duration.

- **Range:** 60 through 2147483647 seconds

infinite Request that the lease never expire.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

DHCP Client

DHCPv6 Client

line-rate (Interfaces)

IN THIS SECTION

- Syntax | 644
- Hierarchy Level | 644
- Description | 645
- Options | 645
- Required Privilege Level | 645
- Release Information | 645

Syntax

```
line-rate
```

Hierarchy Level

```
[edit interfaces interfaces name shdsl-options]
```

Description

Specify a line rate for an G.SHDSL interface.

Options

- `auto`— Automatically selects a line rate.
- `value` — Select the values between 192 kbps and 22784 kbps for the speed of transmission of data on the G.SHDSL connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.0.

link-speed (Interfaces)

IN THIS SECTION

- Syntax | **646**
- Hierarchy Level | **646**
- Description | **646**
- Options | **646**
- Required Privilege Level | **646**

- Release Information | 647

Syntax

```
link-speed                    speed;
```

Hierarchy Level

```
[edit interfaces                    interface-name                    redundant-ether-
options]
```

Description

For redundant Ethernet interfaces in a chassis cluster only, set the required link speed.

Options

speed —For redundant 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).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Release 9.0 of Junos OS.

load-interval

IN THIS SECTION

- Syntax | 647
- Hierarchy Level | 647
- Description | 648
- Options | 648
- Required Privilege Level | 648
- Release Information | 648

Syntax

```
load-interval seconds;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN logical interfaces, specify the interval used to calculate the average load on the network. By default, the average interface load is calculated every 60 seconds.

Options

seconds—Number of seconds at which the average load calculation is triggered.

- **Range:** 20 through 180, in 10-second intervals
- **Default:** 60 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

load-threshold

IN THIS SECTION

- Syntax | 649
- Hierarchy Level | 649
- Description | 649
- Options | 649

Required Privilege Level | 650
Release Information | 650

Syntax

```
load-threshold percent;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN logical interfaces, specify the bandwidth threshold percentage used for adding interfaces. Another link is added to the multilink bundle when the load reaches the threshold value you set. Specify a percentage between 0 and 100.

Options

percent—Bandwidth threshold percentage used for adding interfaces. When set to 0, all available channels are dialed.

- **Range:** 0 through 100 seconds
- **Default:** 100 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

IN THIS SECTION

- Syntax | **650**
- Hierarchy Level | **651**
- Description | **651**
- Default | **652**
- Required Privilege Level | **652**
- Release Information | **652**

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]
```

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.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

Statement modified in Junos OS Release 9.2 for the SRX Series.

RELATED DOCUMENTATION

- Configuring Ethernet Loopback Capability
- Understanding Interfaces

loss-priority (CoS Loss Priority)

IN THIS SECTION

- Syntax | 653
- Hierarchy Level | 653
- Description | 653
- Options | 653

Required Privilege Level | 654
Release Information | 654

Syntax

```
loss-priority level code-points [values ];
```

Hierarchy Level

```
[edit class-of-service loss-priority-maps frame-relay-de map-name]
```

Description

Map CoS values to a packet loss priority (PLP). In Junos OS, classifiers associate incoming packets with a forwarding class (FC) and PLP. PLPs allow you to set the priority for dropping packets. Typically, you mark packets exceeding some service level with a high loss priority—that is, a greater likelihood of being dropped.

Options

level can be one of the following:

- `high`—Packet has high loss priority.
- `medium-high`—Packet has medium-high loss priority.
- `medium-low`—Packet has medium-low loss priority.
- `low`—Packet has low loss priority.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

Understanding Packet Loss Priorities

loss-priority (CoS Rewrite Rules)

IN THIS SECTION

- Syntax | 654
- Hierarchy Level | 655
- Description | 655
- Options | 655
- Required Privilege Level | 655
- Release Information | 655

Syntax

```
loss-priority level;
```

Hierarchy Level

```
[edit class-of-service rewrite-rules type rewrite-name forwarding-class class-name]
```

Description

Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and packet loss priority (PLP). The inputs for the map are the forwarding class and the PLP. The output of the map is the code-point alias or bit pattern.

Options

level can be one of the following:

- high—The rewrite rule applies to packets with high loss priority.
- low—The rewrite rule applies to packets with low loss priority.
- medium-high—The rewrite rule applies to packets with medium-high loss priority.
- medium-low—The rewrite rule applies to packets with medium-low loss priority.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Class of Service User Guide (Security Devices)

loss-priority-maps (CoS Interfaces)

IN THIS SECTION

- Syntax | 656
- Hierarchy Level | 656
- Description | 656
- Options | 657
- Required Privilege Level | 657
- Release Information | 657

Syntax

```
loss-priority-maps {
    frame-relay-de (map-name | default);
}
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

Assign the loss priority map to a logical interface.

Options

- `default`—Apply default loss priority map. The default map contains the following:

```
loss-priority low code-point 0;
loss-priority high code-point 1;
```

- *map-name*—Name of loss priority map to be applied.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

loss-priority-maps (CoS)

IN THIS SECTION

- Syntax | 658
- Hierarchy Level | 658
- Description | 658

Required Privilege Level | 658
Release Information | 659

Syntax

```
loss-priority-maps {
    frame-relay-de loss-priority-map-name {
        loss-priority (high  | low  | medium-high  | medium-low) {
            code-points [bit-string];
        }
    }
}
```

Hierarchy Level

```
[edit class-of-service]
```

Description

Map the loss priority of incoming packets based on CoS values.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces

management (PoE)

IN THIS SECTION

- Syntax | 659
- Hierarchy Level | 659
- Description | 660
- Default | 660
- Options | 660
- Required Privilege Level | 660
- Release Information | 660

Syntax

```
management (class | static);
```

Hierarchy Level

```
[edit poe]
```

Description

Designates how the SRX Series device allocates power to the PoE ports.

Default

static

Options

- `static`—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power configured for the port.
- `class`—When a powered device is connected to a PoE port, the power allocated to it is equal to the maximum power for the class as defined by the IEEE 802.3 AF standard.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

maximum-power (PoE)

IN THIS SECTION

- Syntax | 661
- Hierarchy Level | 661
- Description | 661
- Default | 662
- Options | 662
- Required Privilege Level | 662
- Release Information | 662

Syntax

```
maximum-power watts;
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Maximum amount of power that can be supplied to the port.

Default

15.4 W

Options

`Watts`—The maximum number of watts that can be supplied to the port.

`Range` —0 through 15.4

`Default`—15.4 W

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

mdi-mode

IN THIS SECTION

- Syntax | 663
- Hierarchy Level | 663
- Description | 663
- Options | 664
- Required Privilege Level | 664

Release Information | 665

Syntax

```
mdi-mode mode;
```

Hierarchy Level

```
[edit interfaces interface-name ether-options]
[edit interfaces interface-range range ether-options]
```

Description

You must configure media dependent interface (MDI) properties for a 10-Gigabit Ethernet interface on a copper network port of an EX4550 switch to ensure that both sides of the link are compatible.

MDI refers to the IEEE standard for the interface to an unshielded twisted pair (UTP) cable. Twisted-pair Ethernet standards are such that the majority of cables can be wired "straight-through" (pin 1 to pin 1, pin 2 to pin 2 and so on), but others may need to be wired in the "crossover" form (receive to transmit and transmit to receive).

For most ports, the switch can automatically detect the required connection type and can therefore configure the interface appropriately. However, the switch cannot automatically detect whether the connection type of a 10-Gigabit Ethenet interface on a copper network port is straight-through or crossover.

Therefore, you must set the MDI properties of the local interface of a 10-Gigabit Ethernet interface on a copper network port to ensure that it will work correctly with the other side of the link. When you set this configuration on an interface, you must also disable *auto-negotiation* and set the *speed* to 100m.

NOTE: This configuration does not apply to management Ethernet or console interfaces and it does not apply to 1-Gigabit copper ports.

The proper setting depends both on the type of cable and the setting that is being used on the other side of the link:

- For crossover cables—Set the polarity to match the other side of the link. Specify `mdi` for the switch interface if `mdi` is being used on the other side of the link; specify `mdix` if `mdix` is being used on the other side of the link.
- For straight cables—Set the polarity to be the opposite of the other side link. Specify `mdi` for the switch interface if `mdix` is being used on the other side of the link; specify `mdix` if `mdi` is being used on the other side of the link.

Options

One of the following modes:

`auto` Set the MDI properties to automatic. This setting should *not* be used with 10-Gigabit Ethernet interfaces on a copper network port of an EX4550 switch.

`mdi` Set the MDI properties of the interface to straight through mode. The selection of the mode depends on whether crossover or straight cables are being used and on the setting used on the other side of the link.

`mdix` Set the MDI properties of the interface to crossover mode. The selection of the mode depends on whether crossover or straight cables are being used and on the setting used on the other side of the link.

`force` For SRX Series devices this option enables the MDI properties to auto-mdix always.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.2.

Support for SRX Series devices introduced in Junos OS Release 15.1x49.D80.

RELATED DOCUMENTATION

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- *Interfaces Overview for Switches*
- Junos OS Ethernet Interfaces Configuration Guide
- Interfaces User Guide for Security Devices

media-type (Interfaces)

IN THIS SECTION

- Syntax | 665
- Hierarchy Level | 666
- Description | 666
- Options | 666
- Required Privilege Level | 666
- Release Information | 666

Syntax

```
media-type
```

Hierarchy Level

```
[edit interfaces interface-name media-type]
```

Description

Configure the operating modes for the 2-Port 10 Gigabit Ethernet XPIM.

Options

- copper
- fiber

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Interfaces | 2

minimum-links (Interfaces)

IN THIS SECTION

- Syntax | 667
- Hierarchy Level | 667
- Description | 667
- Options | 668
- Required Privilege Level | 668
- Release Information | 668

Syntax

```
minimum-links number;
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces configured as 802.3ad redundant Ethernet interface link aggregation groups (LAGs) in a chassis cluster only, set the required minimum number of physical child links on the primary node that must be working to prevent the interface from being down. Interfaces configured as redundant Ethernet interface LAGs typically have between 4 and 16 physical interfaces, but only half, those on the primary node, are relevant to the minimum-links setting.

If the number of operating interfaces on the primary node falls below the configured value, it will cause the interface to be down even if some of the interfaces are still working.

For an aggregated ethernet interface, you cannot configure all three configuration options, `bfd-liveness-detection`, `minimum-links`, and `sync-reset` at the same time.

Options

number—For redundant Ethernet interface link aggregation group links, specify the number of physical child links on the primary node in the redundant Ethernet interface that must be working. The default **minimum-links** value is 1. The maximum value is half of the total number of physical child interfaces bound to the redundant Ethernet interface being configured or 8, whichever is smaller.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Release 10.1 of Junos OS.

modem-options

IN THIS SECTION

- Syntax | **669**
- Hierarchy Level | **669**
- Description | **669**
- Required Privilege Level | **669**

- Release Information | 669

Syntax

```
modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
}
```

Hierarchy Level

```
[edit interfaces umd0]
```

Description

For J Series Services Routers, configure a USB port to act as a USB modem.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.2.

mtu

IN THIS SECTION

- Syntax | 670
- Hierarchy Level | 670
- Description | 670
- Options | 671
- Required Privilege Level | 671
- Release Information | 671

Syntax

```
mtu bytes;
```

Hierarchy Level

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

Description

Maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values.

Options

bytes—MTU size.

- **Range:** 0 through 5012 bytes
- **Default:** 1500 bytes (`inet`, `inet6`, and `iso` families), 1448 bytes (`mpls`)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

> *Configuring MRRU on Multilink and Link Services Logical Interfaces*
>
> Junos OS Network Interfaces Library for Routing Devices

native-vlan-id

IN THIS SECTION

- Syntax | **672**
- Hierarchy Level (QFX Series and EX4600) | **672**
- Hierarchy Level (ACX Series, EX Series, SRX Series, M Series, MX Series, and T Series) | **672**
- Hierarchy Level (SRX Series) | **673**

- Description | 673
- Default | 674
- Options | 674
- Required Privilege Level | 674
- Release Information | 674

Syntax

```
native-vlan-id vlan-id;
```

Hierarchy Level (QFX Series and EX4600)

For platforms without ELS:

```
[edit interfaces (QFX Series) interface-name unit 0 family ethernet-switching]
```

For platforms with ELS:

```
[edit interfaces (QFX Series) interface-name]
```

Hierarchy Level (ACX Series, EX Series, SRX Series, M Series, MX Series, and T Series)

```
[edit interfaces ge-fpc/pic/port],
[edit interfaces interface-name]
```

Hierarchy Level (SRX Series)

```
[edit interfaces interface-name ]
```

Description

Configure the VLAN identifier to associate with untagged packets received on the physical interface of a trunk mode interface for the following:

- QFX Series and EX4600
- M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging
- MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging
- T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP
- EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces

The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface, otherwise the untagged packets are dropped. To configure the logical interface, include the `vlan-id` statement (matching the `native-vlan-id` statement on the physical interface) at the `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*`] hierarchy level`.

When the `native-vlan-id` statement is included with the *flexible-vlan-tagging* statement, untagged packets are accepted on the same mixed VLAN-tagged port and on the interfaces that are configured for Q-in-Q tunneling.

When the `native-vlan-id` statement is combined with the *interface-mode* statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.

To configure the logical interface, include the `vlan-id` statement (matching the `native-vlan-id` statement on the physical interface) at the `[edit interfaces` *`interface-name`* `unit` *`logical-unit-number`*`] hierarchy level`.

NOTE: Starting in Junos OS Release 17.1R1, you can send untagged traffic without a native VLAN ID to the remote end of the network. To do this, remove the native VLAN ID from the untagged traffic configuration by setting the `no-native-vlan-insert` statement. If you do not configure this statement, the native VLAN ID is added to the untagged traffic.

Default

By default, the untagged packets are dropped. That is, if you do not configure the `native-vlan-id` option, the untagged packets are dropped.

Options

vlan-id—Numeric identifier of the VLAN.

- **Range:** 1 through 4094

number—VLAN ID number.

- **Range:** (ACX Series routers, SRX Series devices and EX Series switches) 0 through 4094.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.

RELATED DOCUMENTATION

- Configuring Gigabit Ethernet Interfaces (CLI Procedure)
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure)
- *Understanding Bridging and VLANs on Switches*
- *Enabling VLAN Tagging*
- *Configuring Access Mode on a Logical Interface*
- *Configuring the Native VLAN Identifier on Switches With ELS Support*
- Understanding Interfaces | 2
- *Understanding Q-in-Q Tunneling and VLAN Translation*
- *no-native-vlan-insert*
- *Sending Untagged Traffic Without VLAN ID to Remote End*
- *show ethernet-switching interfaces*
- *show vlans*
- *flexible-vlan-tagging*
- Junos OS Network Interfaces Configuration Guide

next-hop-tunnel

IN THIS SECTION

- Syntax | 676
- Hierarchy Level | 676
- Description | 676
- Options | 676
- Required Privilege Level | 676
- Release Information | 677

Syntax

```
next-hop-tunnel gateway-address ipsec-vpn vpn-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family-name]
```

Description

For the secure tunnel (st) interface, create entries in the Next-Hop Tunnel Binding (NHTB) table, which is used to map the next-hop gateway IP address to a particular IP Security (IPsec) Virtual Private Network (VPN) tunnel. NHTB allows the binding of multiple IPsec VPN tunnels to a single IPsec tunnel interface.

Options

- *gateway-address*—Next-hop gateway IP address.
- ipsec-vpn *vpn-name* —VPN to which the next-hop gateway IP address is mapped.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces | 2

no-dns-propagation

IN THIS SECTION

- Syntax | 677
- Hierarchy Level | 677
- Description | 678
- Required Privilege Level | 678
- Release Information | 678

Syntax

```
no-dns-propagation;
```

Hierarchy Level

```
[edit interface interface-name unit unit-number family inet | inet6 dhcp-client]
```

Description

Disable the propagation of DNS information to the kernel.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1X47-D35.

RELATED DOCUMENTATION

Understanding Interfaces | 2

option-refresh-rate (Services)

IN THIS SECTION

- Syntax | 679
- Hierarchy Level | 679
- Description | 679
- Options | 679
- Required Privilege Level | 679
- Release Information | 679

Syntax

```
option-refresh-rate
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specify the option refresh rate.

Options

- packets—Specify the number of packets. The range is from 1 through 480,000.
- seconds—Specify the number of seconds. The range is from 10 through 600.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

pic-mode (Chassis T1 Mode)

IN THIS SECTION

- Syntax | 680
- Hierarchy Level | 680
- Description | 680
- Options | 681
- Required Privilege Level | 681
- Release Information | 681

Syntax

```
pic-mode (clear-channel);
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number ethernet]
```

Description

Configure normal T1 mode or channelized T1 mode.

Options

- `clear-channel`—(default) Normal T1 mode.
- `ct1`—Channelized T1 mode.

> **NOTE**: When chassis clustering is enabled, it is necessary to indicate in the command which node is being configured. In such circumstances, the `edit chassis fpc command` becomes `edit chassis node` *`node-id`* `fpc`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement added in Junos OS Release 10.2.

RELATED DOCUMENTATION

Understanding Interfaces | 2

periodic (Interfaces)

IN THIS SECTION

- Syntax | 682
- Hierarchy Level | 682

- Description | 682
- Options | 682
- Required Privilege Level | 683
- Release Information | 683

Syntax

```
periodic (fast | slow);
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options lacp]
```

Description

For redundant Ethernet interfaces in a chassis cluster only, configure the interval at which the interfaces on the remote side of the link transmit link aggregation control protocol data units (PDUs) by configuring the `periodic` statement on the interfaces on the local side. It is the configuration on the local side that specifies the behavior of the remote side. That is, the remote side transmits link aggregation control PDUs at the specified interval.

Options

- `fast`—Transmit link aggregation control PDUs every second.
- `slow`—Transmit link aggregation control PDUs every 30 seconds.
- **Default:** `fast`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

pool

IN THIS SECTION

- Syntax | 683
- Hierarchy Level | 684
- Description | 684
- Options | 684
- Required Privilege Level | 684
- Release Information | 684

Syntax

```
pool pool-name <priority priority>;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port dialer-options],
[edit interfaces umd0 dialer-options],
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers, for logical and physical ISDN interfaces, specify the dial pool. The dial pool allows logical (dialer) and physical (br-*pim*/0/*port*) interfaces to be bound together dynamically on a per-call basis. On a dialer interface, pool directs the dialer interface which dial pool to use. On br-*pim*/0/*port* interface, pool defines the pool to which the interface belongs.

Options

pool-name—Pool identifier.

priority *priority*—(Physical br-*pim*/0/*port* interfaces only) Specify a priority value of 0 (lowest) to 255 (highest) for the interface within the pool.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

ppp-over-ether

IN THIS SECTION

- Syntax | 685
- Hierarchy Level | 685
- Description | 685
- Required Privilege Level | 686
- Release Information | 686

Syntax

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

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number encapsulation]
```

Description

This encapsulation is used for underlying interfaces of pp0 interfaces. This encapsulation is supported on Fast Ethernet interface, Gigabit Ethernet interface, and Redundant Ethernet interface. When Redundant Ethernet interface is used as underlying interface, an existing pppoe session can be continued in case of failover.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 11.2.

This encapsulation is supported for Redundant Ethernet interface in Junos OS Release 11.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

pppoe

IN THIS SECTION

- Syntax | 687
- Hierarchy Level | 687
- Description | 687
- Options | 687
- Required Privilege Level | 688
- Release Information | 688

Syntax

```
pppoe {
    command binary-file-path;
    disable;
    failover (alternate-media | other-routing-engine);
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Enable users to connect to a network of hosts over a bridge or access concentrator.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the Point-to-Point Protocol over Ethernet process.
- `failover`—Configure the device to reboot if the software process fails four times within 30 seconds, and specify the software to use during the reboot.
  - `alternate-media`—Configure the device to switch to backup media that contains a version of the system if a software process fails repeatedly.
  - `other-routing-engine`—Instruct the secondary Routing Engine to take primary role if a software process fails. If this statement is configured for a process, and that process fails four times within 30 seconds, then the device reboots from the secondary Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

pppoe-options

IN THIS SECTION

- Syntax | 688
- Hierarchy Level | 689
- Description | 689
- Options | 689
- Required Privilege Level | 690
- Release Information | 690

Syntax

```
pppoe-options {
    access-concentrator name ;
    auto-reconnect seconds;
```

```
    (client | server);
    ignore-eol-tag;
    service-name name;
    underlying-interface interface-name;
}
```

Hierarchy Level

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

Description

Configure PPP over Ethernet-specific interface properties.

Options

access-concentrator *name* (SRX Series devices with Point-to-Point Protocol over Ethernet (PPPoE) interfaces) Configure the name of the access concentrator. If you configure a specific access concentrator name on the client and the same access concentrator name server is available, then a PPPoE session is established. If there is a mismatch between the access concentrator names of the client and the server, the PPPoE session gets closed.

auto-reconnect *seconds* Configure the amount of time to wait before reconnecting after a session has terminated.

client Configure the device to operate in the PPPoE client mode.

idle-timeout *seconds* Configure the maximum time that a session can be idle.

ignore-eol-tag Disable the `End-of-List` tag to process the tags after the `End-of-List` tag in a PPPoE Active Discovery Offer (PADO) packet.

service-name *name*	Configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.
server	Configure the device to operate in the PPPoE server mode.
underlying-interface *interface-name*	Configure the interface on which PPP over Ethernet is running.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 12.3X48 to include `ignore-eol-tag` statement.

RELATED DOCUMENTATION

Example: Configuring PPPoE Interfaces | 312

priority (PoE)

IN THIS SECTION

- Syntax | 691
- Hierarchy Level | 691
- Description | 691
- Default | 691

- Options | 691
- Required Privilege Level | 692
- Release Information | 692

Syntax

```
priority (high | low);
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Sets the priority of individual ports. When it is not possible to maintain power to all connected ports, lower-priority ports are powered off before higher priority ports. When a new device is connected on a higher-priority port, a lower-priority port will be powered off automatically if available power is insufficient to power on the higher-priority port. Note that for ports with the same priority configuration, ports on the left are given higher priority than the ports on the right.

Default

low

Options

value—high or low:

- `high`—Specify that this port is to be treated as high priority in terms of power allocation
- `low`—Specify that this port is to be treated as low priority in terms of power allocation.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

profile (Access)

IN THIS SECTION

- Syntax | **692**
- Hierarchy Level | **694**
- Description | **694**
- Options | **695**
- Required Privilege Level | **696**
- Release Information | **696**

Syntax

```
profile profile-name  {
    accounting {
        accounting-stop-on-access-deny;
```

```
        accounting-stop-on-failure;
        duplication;
        duplication-attribute-format;
        duplication-filter;
        duplication-vrf;
        order [accounting-method];
        statistics (time | volume-time);
    }
    address-assignment {
        inet6-pool inet6-pool-name;
        pool pool-name;
    }
    authentication-order (ldap | none | password | radius | s6a | securid);
    charging-service-list;
    client client-name  {
        chap-secret  chap-secret;
        client-group [ group-names ];
        firewall-user {
            password password;
        }
        no-rfc2486;
        pap-password pap-password;
        x-auth ip-address;
    }
    client-name-filter {
        count number;
        domain-name domain-name;
        separator special-character;
    }
    domain-name-server name;
    domain-name-server-inet name;
    domain-name-server-inet6 name;
    jsrc;
    ldap-options {
        assemble {
            common-name common-name;
        }
        base-distinguished-name base-distinguished-name;
        revert-interval seconds;
        search {
            admin-search {
                distinguished-name distinguished-name;
                password password;
```

```
            }
            search-filter search-filter-name;
        }
    }
    ldap-server server-address {
        port port-number;
        retry attempts;
        routing-instance routing-instance-name;
        source-address source-address;
        timeout seconds;
    }
    provisioning-order (gx-plus | jsrc);
    radius;
    radius-options;
    radius-server;
    session-limit-per-username;
    session-options {
        client-group [group-name];
        client-idle-timeout minutes;
        client-session-timeout minutes;
    }
    subscriber;
    wins-server;
}
```

Hierarchy Level

```
[edit access]
```

Description

Create a profile containing a set of attributes that define device management access.

Options

name Profile name

accounting Specifies the accounting options

address-assignment Specify the address assignment pool

authentication-order Order in which authentication mechanisms are used

- Values:
  - ldap—Light weight directory access protocol
  - none—No authentication performed
  - password—Locally configured password in access profile
  - radius—Remote authentication dial-in user service
  - s6a—S6a authentication
  - securid—RSA secure ID authentication

charging-service-list List of used 3gpp charging services

- Values:
  - ocs—Online charging service

client Entity requesting access

client-name-filter Restrictions on client names authenticated on this server

domain-name-server Default DNS server's IPv4 address

domain-name-server-inet DNS server's IPv4 address

domain-name-server-inet6 DNS server's IPv6 address

jsrc Set of JSRC configurations

ldap-options Light weight directory access protocol options

ldap-server Light weight directory access protocol server

preauthentication-order Order in which pre authentication mechanisms are used

- Values:
  - radius—Remote Authentication Dial-In User Service

radius	Set of RADIUS configurations
radius-options	RADIUS options
radius-server	RADIUS server configuration
session-limit-per-username	Maximum number of sessions allowed per username • **Range:** 1 through 16
session-options	Options for an authenticated client's session
subscriber	Locally authenticated subscriber configuration
wins-server	Default WINS server's IPv4 address

Required Privilege Level

access—To view this statement in the configuration.

access-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

`inet6-pool` option is introduced in Junos OS Release 20.3R1.

`none` option is introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

Understanding Interfaces | 2

Understanding User Authentication for Security Devices

Ethernet Switching and Layer 2 Transparent Mode Overview

profiles

IN THIS SECTION

- Syntax | 697
- Hierarchy Level | 697
- Description | 698
- Options | 698
- Required Privilege Level | 698
- Release Information | 698

Syntax

```
profiles {
    profile-name {
        sip-user-id simple-ip-user-id;
        sip-password simple-ip-password;
        access-point-name apn;
        authentication-method (pap | chap | none);
    }
}
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options]
```

Description

Configure a profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network. You can configure up to 16 profiles.

Options

profile-name—Name of the profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

promiscuous-mode (Interfaces)

IN THIS SECTION

Syntax | 699

- Hierarchy Level | 699
- Description | 699
- Required Privilege Level | 699
- Release Information | 700

Syntax

```
promiscuous-mode;
```

Hierarchy Level

```
[edit interfaces interface-name ]
```

Description

Enable promiscuous mode on Layer 3 Ethernet interfaces. When promiscuous mode is enabled on an interface, all packets received on the interface are sent to the central point or Services Processing Unit regardless of the destination MAC address of the packet.

You can also enable promiscuous mode on chassis cluster redundant Ethernet interfaces and on aggregated Ethernet interfaces. If you enable promiscuous mode on a redundant Ethernet interface, promiscuous mode is then enabled on any child physical interfaces. If you enable promiscuous mode on an aggregated Ethernet interface, promiscuous mode is then enabled on all member interfaces.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)

quality (Interfaces)

IN THIS SECTION

- Syntax | 700
- Hierarchy Level | 700
- Description | 701
- Required Privilege Level | 701
- Release Information | 701

Syntax

```
quality <value>;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number radio-router ]
```

Description

This option controls relative link quality weight (value 0–100).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

r2cp

IN THIS SECTION

- Syntax | **702**
- Hierarchy Level | **702**
- Description | **702**
- Options | **702**
- Required Privilege Level | **702**
- Release Information | **703**

Syntax

```
r2cp {
    command binary-file-path;
    disable;
}
```

Hierarchy Level

```
[edit system processes]
```

Description

Specify the Radio-to-Router Control Protocol (R2CP) used to exchange dynamic metric changes in the network that routers use to update the OSPF topologies.

Options

- `command` *`binary-file-path`*—Path to the binary process.
- `disable`—Disable the Radio-to-Router Control Protocol process.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

radio-router (Interfaces)

IN THIS SECTION

- Syntax | 703
- Hierarchy Level | 704
- Description | 704
- Options | 704
- Required Privilege Level | 704
- Release Information | 704

Syntax

```
radio-router {
    bandwidth number;
    credit {
        interval number;
    }
    data-rate number;
    latency number;
    quality number;
    resource number;
```

```
    threshold number;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Point-to-Point Protocol over Ethernet (PPPoE)-based radio-to-router protocols include messages that define how an external system will provide the device with timely information about the quality of a link's connection. They also include a flow control mechanism to indicate how much data the device can forward. The device can then use the information provided in the PPPoE messages to dynamically adjust the interface speed of PPP links.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

redial-delay

IN THIS SECTION

- Syntax | 705
- Hierarchy Level | 705
- Description | 706
- Options | 706
- Required Privilege Level | 706
- Release Information | 706

Syntax

```
redial-delay time;
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options],
[edit logical-systems logical-system-name interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with interfaces configured for ISDN with dialout, specify the delay (in seconds) between two successive calls made by the dialer. To configure callback mode, include the `callback` statement at the `[edit interfaces dl`*n*` unit `*logical-unit-number*` dialer-options]` hierarchy level.

If the `callback` statement is configured, you cannot use the `caller `*caller-id* statement at the `[edit interfaces dl`*n*` unit `*logical-unit-number*` dialer-options]` hierarchy level.

Options

time—Delay (in seconds) between two successive calls.

- **Range:** 2 through 255 seconds
- **Default:** 3 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

ISDN Interfaces Overview

redundancy-group (Interfaces)

IN THIS SECTION

- Syntax | 707
- Hierarchy Level | 707
- Description | 707
- Options | 708
- Required Privilege Level | 708
- Release Information | 708

Syntax

```
redundancy-group                    number          ;
```

Hierarchy Level

```
[edit interfaces                    interface-name                  redundant-ether-options]
```

Description

Specify the redundancy group that a redundant Ethernet interface belongs to.

Options

number —Number of the redundancy group that the redundant interface belongs to. Failover properties of the interface are inherited from the redundancy group.

- **Range:** 1 through 255

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Interfaces User Guide for Security Devices

redundant-ether-options

IN THIS SECTION

- Syntax | **709**
- Hierarchy Level | **709**
- Description | **709**
- Options | **709**
- Required Privilege Level | **710**
- Release Information | **710**

Syntax

```
redundant-ether-options {
     (flow-control | no-flow-control);
    lacp {
        (active | passive);
        periodic (fast | slow);
    }
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
    redundancy-group number;
    source-address-filter mac-address;
    (source-filtering | no-source-filtering);
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure Ethernet redundancy options for a chassis cluster.

In a chassis cluster setup, a redundant Ethernet interface is a pseudointerface that includes at minimum one physical interface from each node of the cluster.

A reth is a special type of interface that has the characteristics of aggregated Ethernet interface.

Options

flow-control Enable flow control.

link-speed Link speed of individual interface that joins the reth interface.

- Values:
  - `100m`—Links are 100 Mbps
  - `10g`—Links are 10 Gbps
  - `10m`—Links are 10 Mbps
  - `1g`—Links are 1Gbps

loopback Enable loopback.

minimum-links Minimum number of active links.

- **Default:** `1`
- **Range:** `1-8`

no-flow-control Do not enable flow control.

no-loopback Do not enable loopback.

no-source-filtering Do not enable source address filtering.

redundancy-group Redundancy group of this interface.

- **Range:** `1-128`

source-filtering Enable source address filtering.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Example: Enabling Eight-Queue Class of Service on Redundant Ethernet Interfaces on SRX Series Devices in a Chassis Cluster

Example: Configuring Chassis Cluster Redundant Ethernet Interfaces

redundant-parent (Interfaces Fast Ethernet)

IN THIS SECTION

- Syntax | 711
- Hierarchy Level | 711
- Description | 711
- Options | 712
- Required Privilege Level | 712
- Release Information | 712

Syntax

```
redundant-parent interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name fastether-options]
```

Description

Configure Fast Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.

Options

interface —Parent redundant interface of the Fast Ethernet interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

redundant-parent (Interfaces Gigabit Ethernet)

IN THIS SECTION

- Syntax | 713
- Hierarchy Level | 713
- Description | 713
- Options | 713
- Required Privilege Level | 713
- Release Information | 713

Syntax

```
redundant-parent interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

Description

Configure Gigabit Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.

Options

interface —Parent redundant interface of the Gigabit Ethernet interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 9.0 of Junos OS.

request pppoe connect

IN THIS SECTION

- Syntax | 714
- Description | 714
- Options | 714
- Required Privilege Level | 714
- Output Fields | 715
- Sample Output | 715
- Release Information | 715

Syntax

```
request pppoe connect
```

Description

Connect all sessions that are down.

Options

`pppoe interface name`— (Optional) Connect to a specified session.

Required Privilege Level

maintenance

Output Fields

When you enter this command, this command returns no output.

Sample Output

request pppoe connect

```
user@host> request pppoe connect
```

Release Information

Statement supported on SRX300, SRX320, SRX340, and SRX345 is introduced in Junos OS Release 15.1X49-D60.

Statement supported on SRX1500 and vSRX instances is introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

- Understanding PPPoE Interfaces
- Example: Configuring PPPoE Interfaces

request pppoe disconnect

IN THIS SECTION

- Syntax | 716
- Description | 716
- Options | 716
- Required Privilege Level | 716

- Output Fields | 716
- Sample Output | 717
- Release Information | 717

Syntax

```
request pppoe disconnect
```

Description

Disconnect all active sessions.

Options

`session id` — (Optional) Disconnect the session for which the session ID is specified.

`pppoe interface name`— (Optional) Disconnect the session for a specific pppoe interface name.

Required Privilege Level

maintenance

Output Fields

When you enter this command, this command returns no output.

Sample Output

request pppoe disconnect

```
user@host> request pppoe disconnect
```

Release Information

Statement supported on SRX300, SRX320, SRX340, and SRX345 is introduced in Junos OS Release 15.1X49-D60.

Statement supported on SRX1500 and vSRX instances is introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

- Understanding PPPoE Interfaces
- Example: Configuring PPPoE Interfaces

resource (Interfaces)

IN THIS SECTION

- Syntax | 718
- Hierarchy Level | 718
- Description | 718
- Required Privilege Level | 718
- Release Information | 718

Syntax

```
resource number;
```

Hierarchy Level

```
[edit interfaces interface-name radio-router ]
```

Description

This option controls the resource weight (value 1–100).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

retransmission-attempt (DHCP Client)

IN THIS SECTION

- Syntax | 719
- Hierarchy Level | 719
- Description | 719
- Options | 720
- Required Privilege Level | 720
- Release Information | 720

Syntax

```
retransmission-attempt number;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet dhcp]
```

Description

Specify the number of times the device retransmits a Dynamic Host Control Protocol (DHCP) packet if a DHCP server fails to respond. After the specified number of attempts, no further attempts at reaching a server are made.

Options

number Number of retransmit attempts.

Range:

- - For IPv4 — 0 through 50000 from Junos OS Release 17.3R1 onwards and 0 through 6 on Junos OS earlier releases.
  - For IPv6 — 0 through 9.
- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

> *Configuring a DHCP Client*
>
> *interfaces*
>
> unit
>
> family

retransmission-interval (DHCP Client)

IN THIS SECTION

- Syntax | 721
- Hierarchy Level | 721
- Description | 721
- Options | 721
- Required Privilege Level | 722
- Release Information | 722

Syntax

```
retransmission-interval seconds;
```

Hierarchy Level

```
[edit interfaces   interface-name  unit logical-unit-number        family family-name dhcp]
```

Description

Specify the time between successive retransmission attempts.

Options

seconds—Number of seconds between successive retransmission.

- **Range:** 4 through 64 seconds
- **Default:** 4 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Release 8.5 of Junos OS.

roaming-mode

IN THIS SECTION

- Syntax | 722
- Hierarchy Level | 723
- Description | 723
- Options | 723
- Required Privilege Level | 723
- Release Information | 723

Syntax

```
roaming-mode (home-only | automatic)
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options]
```

Description

Specify whether the 3G wireless modem interface can access networks other than the home network.

Options

- `home-only`—No roaming is allowed.
- `automatic`—Allows access to networks other than the home network. This is the default.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

scheduler-map (CoS Virtual Channels)

IN THIS SECTION

- Syntax | 724
- Hierarchy Level | 724
- Description | 724
- Options | 724
- Required Privilege Level | 725
- Release Information | 725

Syntax

```
scheduler-map map-name;
```

Hierarchy Level

```
[edit class-of-service virtual-channel-groups group-name virtual-channel-name]
```

Description

Apply a scheduler map to this virtual channel.

Options

map-name—Name of the scheduler map.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

- *default (CoS)*
- *shaping-rate (CoS Virtual Channels)*
- *virtual-channel-group (CoS Interfaces)*
- *virtual-channel-groups*
- *virtual-channels*

select-profile

IN THIS SECTION

- Syntax | **726**
- Hierarchy Level | **726**
- Description | **726**
- Options | **726**
- Required Privilege Level | **726**
- Release Information | **726**

Syntax

```
select-profile profile-name
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options]
```

Description

Select the active profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network.

Options

profile-name—Name of a configured profile that is to be used to establish a data call.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

server-address

IN THIS SECTION

- Syntax | 727
- Hierarchy Level | 727
- Description | 728
- Default | 728
- Options | 728
- Required Privilege Level | 728
- Release Information | 728

Syntax

```
server-address  ip-address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet dhcp]
```

Description

Specify the address of the DHCP server that the client should accept DHCP offers from. If this option is included in the DHCP configuration, the client accepts offers only from this server and ignores all other offers.

Default

The client accepts the first offer it receives from any DHCP server.

Options

ip-address	DHCP server address.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- *Configuring a DHCP Client*
- *interfaces*
- unit
- family

shaping-rate (CoS Interfaces)

IN THIS SECTION

- Syntax | 729
- Hierarchy Level | 729
- Description | 729
- Default | 730
- Options | 730
- Required Privilege Level | 730
- Release Information | 731

Syntax

```
shaping-rate rate <overhead bytes> ;
```

Hierarchy Level

```
[edit class-of-service interfaces interface-name],
[edit class-of-service interfaces interface-name unit logical-unit-number]
```

Description

For logical interfaces on which you configure packet scheduling, configure traffic shaping by specifying the amount of bandwidth to be allocated to the logical interface.

Logical and physical interface traffic shaping can be configured together. This means you can include the `shaping-rate` statement at the `[edit class-of-service interfaces interface interface-name]` hierarchy level *and* the `[edit class-of-service interfaces interface interface-name unit logical-unit-number]` hierarchy level. If you

configure traffic shaping at both the logical and physical interface levels, the logical interface shaping credit is checked and updated before the physical interface shaping credit.

Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the `shaping-rate` statement at the `[edit class-of-service traffic-control-profiles]` hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.

On the physical interface, you can set the Layer 2 overhead adjustment to the shaping rate calculation at egress.

Default

If you do not include this statement at the `[edit class-of-service interfaces interface` *`interface-name`* `unit` *`logical-unit-number`*`]` hierarchy level, the default logical interface bandwidth is the average of unused bandwidth for the number of logical interfaces that require default bandwidth treatment. If you do not include this statement at the `[edit class-of-service interfaces interface` *`interface-name`*`]` hierarchy level, the default physical interface bandwidth is the average of unused bandwidth for the number of physical interfaces that require default bandwidth treatment.

Options

rate Peak rate, in bits per second (bps). 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).

- **Range:** 1000 through 6,400,000,000,000 bps

overhead Layer 2 shaping overhead adjustment to be applied at egress (bytes).

- **Range:** -62 through 192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

`overhead` option introduced in Junos OS Release 18.1.

RELATED DOCUMENTATION

policer-overhead

simple-filter (Interfaces)

IN THIS SECTION

- Syntax | 731
- Hierarchy Level | 731
- Description | 732
- Options | 732
- Required Privilege Level | 732
- Release Information | 732

Syntax

```
simple-filter;
```

Hierarchy Level

```
[edit interfaces interfaces-name unit logical-unit-number family family-name]
```

Description

Apply a simple filter to an interface. You can apply simple filters on ingress interfaces only.

Options

input *filter-name*: Name of one filter to evaluate when packets are received on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

sip-password

IN THIS SECTION

- Syntax | 733
- Hierarchy Level | 733
- Description | 733

- Options | 733
- Required Privilege Level | 733
- Release Information | 734

Syntax

```
sip-password simple-ip-password;
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Configure the password provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

simple-ip-password—Password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

sip-user-id

IN THIS SECTION

- Syntax | 734
- Hierarchy Level | 734
- Description | 735
- Options | 735
- Required Privilege Level | 735
- Release Information | 735

Syntax

```
sip-user-id simple-ip-user-id;
```

Hierarchy Level

```
[edit interfaces interface-name cellular-options gsm-options profiles profile-name]
```

Description

Configure the username provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.

Options

simple-ip-user-id—Username.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

source-address-filter (Interfaces)

IN THIS SECTION

- Syntax | 736
- Hierarchy Level | 736
- Description | 736
- Options | 736
- Required Privilege Level | 737
- Release Information | 737

Syntax

```
source-address-filter mac-address;
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces, specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the `source-filtering` statement in the configuration to enable source address filtering.

Be sure to update the MAC address if the remote Ethernet card is replaced. Replacing the interface card changes the MAC address. Otherwise, the interface cannot receive packets from the new card.

NOTE:

- Software based MAC limiting is supported on SRX300, SRX320, and SRX340 devices.

  A maximum of 32 devices are supported per device.

Options

mac-address —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* (for example, `00:11:22:33:44:55`) or *nnnn:nnnn:nnnn* (for example, `0011.2233.4455`). You can configure up to 64 source addresses. To specify more than one address, include multiple *mac-address* options in the `source-address-filter` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

source-filtering (Interfaces)

IN THIS SECTION

- Syntax | 737
- Hierarchy Level | 738
- Description | 738
- Required Privilege Level | 738
- Release Information | 738

Syntax

```
(source-filtering | no-source-filtering);
```

Hierarchy Level

```
[edit interfaces interface-name redundant-ether-options]
```

Description

For redundant Ethernet interfaces, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the `source-address-filter` statement.

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

By default, source address filtering is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement modified in Junos OS Release 9.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

speed (Interfaces)

IN THIS SECTION

- Syntax | 739
- Hierarchy Level | 739
- Description | 739
- Options | 739
- Required Privilege Level | 740
- Release Information | 740

Syntax

```
speed (100m |10m | 1g);
```

Hierarchy Level

```
[edit interfaces interface-name speed]
```

Description

Configure the operating speed for the 2-Port 10 Gigabit Ethernet XPIM.

Options

- 100m — Link speed of 100 Mbps

- 10g — Link speed of 10 Gbps
- 10m — Link speed of 10 Mbps
- 1g — Link speed of 1 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface | 222

speed (Gigabit Ethernet interface)

IN THIS SECTION

- Syntax | 741
- Hierarchy Level | 741
- Description | 741
- Options | 741
- Required Privilege Level | 742
- Release Information | 742

Syntax

```
speed (1g |10g);
```

Hierarchy Level

```
[edit interfaces interface-name gigether-options]
```

Description

Configure the operating speed of the 8-port 10-Gigabit Ethernet PIC from default 10-Gbps port speed to 1-Gbps port speed. Each of the interfaces in the 8-port 10-Gigabit Ethernet PIC can be independently configured to 1Gbps or 10Gbps speeds. For information about platforms support, see hardware compatibility tool (HCT).

Autonegotiation is automatically disabled when 1-Gbps speed is configured on the interfaces.

On 1/10-Gbps capable Gigabit Ethernet SFP interfaces, the duplex is always full and the speed matches that of the inserted optic. These interfaces support either 1-Gbps or 10-Gbps SFP optics. For SRX devices, 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 config will always show as xe-. If a speed configuration is changed, you cannot change it again in the next 180 seconds. The interface link might drop down, if you try to change the speed configuration again within 180 seconds of the first speed configuration change. The 8x10-Gbps ports supports multiple port speeds, that is, some ports operates at 10G speed and some at 1G speed. To view the speed configured for the interface, execute the `show interfaces extensive` command. The `Speed Configuration` field's value of `1G` or `AUTO` in the command output indicates whether the current operation speed of the interface is 1 Gbps or the default 10 Gbps, respectively.

Options

- 1g — Link speed of 1 Gbps
- 10g — Link speed of 10 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

speed (Chassis Cluster) | 1112

spid1

IN THIS SECTION

- Syntax | 742
- Hierarchy Level | 743
- Description | 743
- Options | 743
- Required Privilege Level | 743
- Release Information | 743

Syntax

```
spid1 spid1-string;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

Configure the Service Profile Identifier (SPID).

Options

spid1-string—Numeric SPID.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

spid2

IN THIS SECTION

Syntax | 744

- Hierarchy Level | 744
- Description | 744
- Options | 744
- Required Privilege Level | 744
- Release Information | 745

Syntax

```
spid2 spid2-string;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

Configure an additional SPID.

Options

spid2-string—Numeric SPID.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

static-tei-val

IN THIS SECTION

- Syntax | 745
- Hierarchy Level | 745
- Description | 746
- Options | 746
- Required Privilege Level | 746
- Release Information | 746

Syntax

```
static-tei-val value;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

For J Series Services Routers only. Statically configure the Terminal Endpoint Identifier (TEI) value. The TEI value represents any ISDN-capable device attached to an ISDN network that is the terminal endpoint. TEIs are used to distinguish between several different devices using the same ISDN links.

Options

value—Value between 0 through 63.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

switch-type

IN THIS SECTION

- Syntax | **747**
- Hierarchy Level | **747**
- Description | **747**
- Options | **747**
- Required Privilege Level | **747**

- Release Information | 748

Syntax

```
switch-type (att5e | etsi | ni1 | ntdms-100)
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

For J Series Services Routers only. Configure the ISDN variant supported.

Options

`att5e`—AT&T switch variant.

`etsi`—European Telecommunications Standards Institute switch variant.

`ni1`—National ISDN 1 switch variant.

`ntdms-100`—Northern Telecom DMS-100.

`ntt`—NTT Group switch for Japan.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

t310

IN THIS SECTION

- Syntax | 748
- Hierarchy Level | 748
- Description | 749
- Options | 749
- Required Privilege Level | 749
- Release Information | 749

Syntax

```
t310-value seconds;
```

Hierarchy Level

```
[edit interfaces br-pim/0/port isdn-options]
```

Description

For ISDN interfaces, configure the Q.931-specific timer for T310, in seconds. The Q.931 protocol is involved in the setup and termination of connections.

Options

seconds—Timer value, in seconds.

- **Range:** 1 through 65,536 seconds
- **Default:** 10 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

tei-option

IN THIS SECTION

- Syntax | **750**
- Hierarchy Level | **750**
- Description | **750**
- Options | **750**

Required Privilege Level | 750
Release Information | 751

Syntax

```
tei-option (first-call | power-up);
```

Hierarchy Level

```
[edit interfaces br-pim/0/portisdn-options ]
```

Description

For ISDN interfaces, configure when the Terminal Endpoint Identifier (TEI) negotiates with the ISDN provider.

Options

`first-call`—Activation does not occur until the call setup is sent.

`power-up`—Activation occurs when the Services Router is powered on.

- **Default:** `power-up`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

telemetries (PoE)

IN THIS SECTION

- Syntax | 751
- Hierarchy Level | 752
- Description | 752
- Default | 752
- Required Privilege Level | 752
- Release Information | 752

Syntax

```
telemetries {
    disable;
    duration hours;
    interval minutes;
}
```

Hierarchy Level

```
[edit poe interface (all | interface-name)]
```

Description

Allow logging of per-port PoE power consumption. The telemetries section must be explicitly specified to enable logging. If left unspecified, telemetries is disabled by default.

Default

If the telemetries statement is specified, logging is enabled with the default values for interval and duration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

template-refresh-rate (Services)

IN THIS SECTION

- Syntax | 753
- Hierarchy Level | 753
- Description | 753
- Options | 753
- Required Privilege Level | 754
- Release Information | 754

Syntax

```
template-refresh-rate;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Specify the template refresh rate.

Options

- packets—Specify the number of packets. The range is from 1 through 480,000.

- seconds—Specify the number of seconds. The range is from 10 through 600.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Understanding Interfaces | 2

threshold (Interfaces)

IN THIS SECTION

- Syntax | 755
- Hierarchy Level | 755
- Description | 755
- Required Privilege Level | 755
- Release Information | 755

Syntax

```
threshold <value>;
```

Hierarchy Level

```
[edit interfaces interface-name radio-router ]
```

Description

This option controls the percentage of bandwidth change required for routing updates.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

traceoptions (Interfaces)

IN THIS SECTION

- Syntax | 756
- Hierarchy Level | 756
- Description | 756
- Options | 757
- Required Privilege Level | 757
- Release Information | 757

Syntax

```
traceoptions
```

Hierarchy Level

```
[edit interfaces interface-name traceoptions]
```

Description

Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.

Options

flag - Tracing parameters

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

PPPoE-Based Radio-to-Router Protocols Overview

update-server

IN THIS SECTION

- Syntax | **758**
- Hierarchy Level | **758**
- Description | **758**
- Required Privilege Level | **758**
- Release Information | **758**

Syntax

```
update-server;
```

Hierarchy Level

```
[edit interfaces interface-name unit  logical-unit-number  family inet dhcp]
```

Description

Propagate TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch, router, or device.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

RELATED DOCUMENTATION

- *Configuring a DHCP Client*
- *Example: Configuring the Device as a DHCP Client*
- *interfaces*
- unit
- family

vbr rate

IN THIS SECTION

- Syntax | 759
- Hierarchy Level | 759
- Description | 759
- Options | 759
- Required Privilege Level | 760
- Release Information | 760

Syntax

```
vbr rate;
```

Hierarchy Level

```
[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
```

Description

For ATM encapsulation only, define a variable bit rate bandwidth utilization in the traffic-shaping profile.

Options

- Burst Size–The maximum burst size that can be sent at the peak rate.

- Peak Rate–The maximum instantaneous rate at which the user will transmit.
- Sustained Rate–The average rate as measured over a long interval.
- CDVT–Cell Delay Variation Tolerance in microseconds (range: 1 – 9999).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces | 2

vdsl-profile

IN THIS SECTION

- Syntax | 761
- Hierarchy Level | 761
- Description | 761
- Options | 761
- Required Privilege Level | 762
- Release Information | 762

Syntax

```
vdsl-profile
```

Hierarchy Level

```
[edit interfaces interface-name vdsl-options]
```

Description

Configure the type of VDSL2 profiles. A profile is a table that contains a list of preconfigured VDSL2 settings.

Options

- Auto (default)
- 8a
- 8b
- 8c
- 8d
- 12a
- 12b
- 17a

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

VDSL2 Interface Support on SRX Series Devices

vendor-id (Interfaces)

IN THIS SECTION

- Syntax | 762
- Hierarchy Level | 763
- Description | 763
- Options | 763
- Required Privilege Level | 763
- Release Information | 763

Syntax

```
vendor-id          vendor-id        ;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number
family family-name dhcp]
```

Description

Configure a vendor class ID for the Dynamic Host Configuration Protocol (DHCP) client.

Options

vendor-id —vendor class ID.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Interfaces | 2

watch-list

IN THIS SECTION

- Syntax | 764
- Hierarchy Level | 764
- Description | 764
- Options | 765
- Required Privilege Level | 765
- Release Information | 765

Syntax

```
watch-list {
    [ routes ];
}
```

Hierarchy Level

```
[edit interfaces dln unit logical-unit-number dialer-options]
```

Description

On J Series Services Routers with ISDN interfaces, configure an ISDN list of routes to watch. Used only for dialer watch.

Options

routes—IP prefix of a route. Specify one or more. The primary interface is considered up if there is at least one valid route for any of the addresses in the watch list to an interface other than the backup interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

web-authentication (Interfaces)

IN THIS SECTION

- Syntax | **766**
- Hierarchy Level | **766**
- Description | **766**
- Options | **766**
- Required Privilege Level | **766**
- Release Information | **767**

Syntax

```
web-authentication {
    http;
    https;
    redirect-to-https;
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family family-name address]
```

Description

Enable the Web authentication process for firewall user authentication.

Options

`http`—Enable HTTP service.

`https`—Enable authentication through HTTPS.

`redirect-to-https`—Redirect Web authentication to HTTPS.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

Support for `https` and `redirect-to-https` introduced for SRX5400, SRX5600, and SRX5800 Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

RELATED DOCUMENTATION

Understanding Interfaces | 2

wlan

IN THIS SECTION

- Syntax (SRX Series) | 767
- Hierarchy Level | 769
- Description | 769
- Options | 769
- Required Privilege Level | 771
- Release Information | 772

Syntax (SRX Series)

```
wlan {
    access-point name {
        description description;
        interface wl interface;
        access-point-options {
            country country-code;
```

```
        }
        location location;
        mac-address mac-address;
        radio (1| 2) {
            radio-options {
            channel {
                number (auto | channel-number);
                bandwidth (20 | 40 | 80);
            }
            mode (g | gn | an | acn);
            radio-off;
            transmit-power percent;
            }
        }
    }
    virtual-access-point id {
        description description;
        no-broadcast-ssid;
        maximum-stations number;
        station-isolation;
        upload-limit upload-limit-rate;
        download-limit download-limit-rate;
        ssid ssid;
        vlan vlan-id;
        station-mac-filter (allow-list | deny-list) {
            mac-address addr1 addr2;
        }
    }
    security {
        none;
        wpa-enterprise {
            cipher-suites ccmp;
            radius-server ip-address;
            radius-port port;
            radius-key secret-key;
            wpa-version v2;
            }
            wpa-personal {
                cipher-suites ccmp;
                key (ascii | hex) key;
                wpa-version v2;
            }
        }
```

```
    }
}
```

Hierarchy Level

```
[edit wlan access-point name]
```

Description

Configure WLAN properties on SRX Series devices.

Options

access-point name	Name of the wireless access point.
interface	Wireless LAN interface (wl-*x*/0/0) created for the access-point setting. To support Mini-PIM in chassis cluster mode, you can configure two WLAN interfaces (wl-*x*/0/0) and (wl-*y*/0/0) on both nodes, then the WLAN configuration for access-point is pushed to the two wireless LAN interface cards on both the nodes.
description	Description of the access point and virtual access point (VAP). The maximum length is 64 characters.
country	The country code.
location	Location of the access point. The maximum number of characters you can use is 64.
channel number (auto | channel number)	Channel number of the radio. If you select `auto`, then the Mini-PIM chooses the channel automatically.
bandwidth	Radio 1 (5 GHz) supports bandwidth of 20MHz, 40MHz, and 80MHz, whereas Radio 2 (2.4 GHz) supports bandwidth of 20MHz and 40MHz. The default value is 20MHz for 2.4GHz and 40MHz for 5GHz.

radio modes (an | acn | gn | g)

Mode for the radio operation.

Radio 1 supports the following modes:

- **an** 802.11a and 802.11n clients operating in 5-GHz frequency can connect to the access point.

- **acn** 802.11a, 802.11b, 802.11n and 802.11ac clients operating in 5-GHz frequency can connect to the access point.

Radio 2 supports the following mode:

- **gn** 802.11g, 802.11b, and 802.11n clients operating in 2.4-GHz frequency can connect to the access point. This is the default mode for this radio.

  g 802.11g clients operating in 2.4-GHz frequency can connect to the access point supported from Junos OS Release 20.4R1.

- **radio-off** Radio is turned off.

- **transmit-power** The percentage of transmit power.

  NOTE: When you configure the transmit power, the Mini-PIM card will fix transmit power to the specified value set, in this case the power by rate functionality does not work. So it is recommended not to set transmit power to a specified value. When you do not configure the transmit power (do not fix the transmit power to a specified value), the power by rate functionality works. If you configure the transmit power percentage to 100, then it chooses the option "auto", the behavior is same as no transmit power configuration is done and power by rate functionality will work.

no-broadcast-ssid Disable broadcast SSID. By default, the broadcast SSID is enabled.

maximum-stations The number of maximum clients that can be connected to the virtual access point. The range is 1 through 127.

station-isolation Isolate the clients connected to the same VAP.

security (none | wpa-enterprise | wpa-personal)	Security settings for the VAP. WPA enterprise is a Wi-Fi Alliance standard that uses RADIUS server authentication with AES-CCMP. This mode allows the use of high-security encryption along with centrally managed user authentication. WPA personal is a Wi-Fi Alliance standard that uses preshared key (PSK) authentication with AES-CCMP . Only WPA2 standards are supported on Wi-Fi Mini-PIM.
none	No security. The data transferred between clients and the access point is not encrypted. This method allows clients to associate with the access point without any authentication.
cipher-suites (ccmp)	Select the appropriate WPA cipher algorithm. The value is CCMP algorithms.
radius-server	IP address of the radius server.
radius-port	Port number of the RADIUSs server. The default value is 1812.
radius-key	Secret key of the RADIUS. The maximum number of characters you can use is 64.
key (ascii | hex)	WPA shared key. The range of key length is 8 through 63 for ASCII or 8 through 64 hexadecimal characters.
wpa-version (v2)	Version of the WPA version. Only supported value is WPA2.
upload-limit	Specify the upload rate limit. The range is from 256 Kbps through 1,048,576 Kbps.
download-limit	Specify the download rate limit. The range is from 256 Kbps through 1,048,576 Kbps.
ssid	SSID value for the virtual access point. The range is 2 through 32. SSID value can include only letters, numbers and five special characters—hyphen (-), underscore (_), at (@), hash (#), and period (.) in the value of the SSID.
vlan-id	VLAN ID for the virtual access point. The range is 1 through 4094. The default value is 1.
station-mac-filter (allow-list| deny-list)	Specify the MAC filter. You can set either allow the mac address list or deny it. The MAC address format is like xx:xx:xx:xx:xx:xx. The maximum number of the MAC addresses listed is 16.

Required Privilege Level

routing—To view this statement in the configuration.

routing–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.

Release Information

Statement introduced in Junos OS Release 19.4R1.

RELATED DOCUMENTATION

Wi-Fi Mini-Physical Interface Module Overview | 443

Configure Wi-Fi Mini-PIM | 446

12

CHAPTER

Operational Commands

clear oam ethernet connectivity-fault-management path-database | 775

clear dhcpv6 server binding (Local Server) | 776

clear ethernet-switching statistics mac-learning | 778

clear interfaces statistics swfabx | 780

clear ipv6 neighbors | 781

clear lacp statistics interfaces | 783

restart | 785

request modem wireless create-profile | 802

request modem wireless fota | 804

request modem wireless sim-lock | 806

request modem wireless sim-unlock | 808

request wlan access-point packet capture | 810

show chassis fpc (View) | 812

show chassis hardware (View) | 824

show ethernet-switching mac-learning-log | 845

show ethernet-switching table | 851

show igmp-snooping route (View) | 885

show interfaces | 887

show interfaces diagnostics optics | 1023

show interfaces flow-statistics | 1030

show interfaces queue | 1037

show interfaces statistics (View) | 1044

show interfaces terse zone | 1046

show ipv6 neighbors | 1047

show lacp interfaces (View) | 1050

show lacp statistics interfaces (View) | 1056

show modem wireless firmware | 1059

show modem wireless network | 1063

show modem wireless profiles | 1068

show oam ethernet link-fault-management | 1071

show poe controller (View) | 1080

show pppoe interfaces | 1082

show pppoe statistics | 1088

show poe telemetries | 1092

show services accounting | 1095

show services accounting aggregation (View) | 1099

show services accounting aggregation template (View) | 1100

show services accounting flow-detail (View) | 1102

show wlan access-points | 1103

speed (Chassis Cluster) | 1112

clear oam ethernet connectivity-fault-management path-database

IN THIS SECTION

- Syntax | 775
- Description | 775
- Options | 775
- Required Privilege Level | 776
- Sample Output | 776
- Release Information | 776

Syntax

```
clear oam ethernet connectivity-fault-management path-database maintenance-domain md-name
maintenance-association ma-name host <mac-addr>
```

Description

Clear the relevant path information from the database for the specified remote host.

Options

host	MAC address of remote host in xx:xx:xx:xx:xx:xx format.
maintenance-association	Name of the maintenance association.
maintenance-domain	Name of the maintenance domain.

Required Privilege Level

clear

Sample Output

clear oam ethernet connectivity-fault- management path-database

```
user@host> clear oam ethernet connectivity-fault-management path-database maintenance-domain
private maintenance-association private-ma 00:00:5E:00:53:AA
Path database entries cleared for the remote-host
```

Release Information

Statement introduced in Junos OS Release 12.1X44-D10.

RELATED DOCUMENTATION

show oam ethernet connectivity-fault-management path-database

clear dhcpv6 server binding (Local Server)

IN THIS SECTION

- Syntax | 777
- Description | 777
- Options | 777
- Required Privilege Level | 778

Release Information | 778

Syntax

```
clear dhcpv6 server binding
<all | client-id | ip-address | session-id>
<interface interface-name>
<routing-instance routing-instance-name>
```

Description

Clear the binding state of a DHCPv6 client from the client table on the DHCPv6 local server.

Options

- all—(Optional) Clear the binding state for all DHCPv6 clients.
- *client-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified client ID (option 1).
- *ip-address*—(Optional) Clear the binding state for the DHCPv6 client with the specified address.
- *session-id*—(Optional) Clear the binding state for the DHCPv6 client with the specified session ID.
- interface *interface-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.
- routing-instance *routing-instance-name*—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

show dhcpv6 server binding (View)

clear ethernet-switching statistics mac-learning

IN THIS SECTION

- Syntax | 778
- Description | 779
- Options | 779
- Required Privilege Level | 779
- Sample Output | 779
- Release Information | 779

Syntax

clear ethernet-switching statistics mac-learning

Description

Clear the media access control (MAC) learning statistics.

Options

- `none`—Clear MAC learning statistics on all interfaces.
- `interface` *`interface-name`*—(Optional) Clear MAC learning statistics on the specified interface.

Required Privilege Level

view

Sample Output

clear ethernet-switching statistics mac-learning

```
user@host> clear ethernet-switching statistics mac-learning
```

clear ethernet-switching statistics mac-learning interface interface-name

```
user@host> clear ethernet-switching statistics mac-learning interface interface-name
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

show ethernet-switching table

clear interfaces statistics swfabx

IN THIS SECTION

- Syntax | 780
- Description | 780
- Required Privilege Level | 780
- Output Fields | 781
- Sample Output | 781
- Release Information | 781

Syntax

```
clear interfaces statistics <swfab0 | swfab1>
```

Description

Clear interface statistics for the specified swfab interface.

Required Privilege Level

clear

Output Fields

When you enter this command, interface statistics for swfab0 and swfab1 are cleared.

Sample Output

clear interfaces statistics <swfab0 | swfab1>

```
user@host> clear interfaces statistics <swfab0 | swfab1>
```

Release Information

Command introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

show interfaces swfabx

clear ipv6 neighbors

IN THIS SECTION

- Syntax | **782**
- Description | **782**
- Options | **782**
- Required Privilege Level | **782**
- Output Fields | **782**
- Sample Output | **783**
- Release Information | **783**

Syntax

```
clear ipv6 neighbors
<all | host hostname>
```

Description

Clear IPv6 neighbor cache information.

Options

none	Clear all IPv6 neighbor cache information.
all	(Optional) Clear all IPv6 neighbor cache information.
host *hostname*	(Optional) Clear the information for the specified IPv6 neighbors.

Required Privilege Level

clear

Output Fields

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
11:11::2            00:19:e2:4b:61:83  deleted
            12:12::2            00:19:e2:4b:61:83  deleted
            10:1::2             00:00:0a:00:00:00  deleted
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

show ipv6 neighbors | 1047

clear lacp statistics interfaces

IN THIS SECTION

- Syntax | 784
- Description | 784
- Options | 784
- Required Privilege Level | 784
- Output Fields | 784
- Release Information | 784

Syntax

```
clear lacp statistics interfaces <interface-name>
```

Description

Clear the LACP statistics. If you do not specify an interface name, LACP statistics for all interfaces are cleared.

Options

interface-name	(Optional) Name of an interface.

Required Privilege Level

clear

Output Fields

This command produces no output.

Release Information

Command modified in Junos OS Release 10.2.

RELATED DOCUMENTATION

show lacp statistics interfaces (View) | 1056

Verifying LACP on Redundant Ethernet Interfaces

restart

IN THIS SECTION

- Syntax | 785
- Syntax (ACX Series Routers) | 786
- Syntax (EX Series Switches) | 786
- Syntax (MX Series Routers) | 787
- Syntax (QFX Series) | 787
- Syntax (Routing Matrix) | 788
- Syntax (SRX Series) | 788
- Syntax (TX Matrix Routers) | 789
- Syntax (TX Matrix Plus Routers) | 789
- Syntax (QFX Series) | 789
- Syntax (Junos OS Evolved) | 790
- Description | 790
- Options | 791
- Required Privilege Level | 800
- Output Fields | 800
- Sample Output | 800
- Release Information | 801

Syntax

```
restart
<adaptive-services |ancpd-service | application-identification |audit-process | auto-
configuration |captive-portal-content-delivery |ce-l2tp-service |chassis-control | class-of-
service |clksyncd-service |database-replication|datapath-trace-service |dhcp-service | diameter-
```

```
service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | ethernet-connectivity-
fault-management |ethernet-link-fault-management |event-processing | firewall | general-
authentication-service | gracefully | iccp-service |idp-policy | immediately |interface-control
| ipsec-key-management | kernel-health-monitoring | kernel-replication | l2-learning | l2cpd-
service | l2tp-service | l2tp-universal-edge | lacp | license-service |link-management |local-
policy-decision-function |mac-validation |mib-process | mountd-service |mpls-traceroute |mspd |
multicast-snooping |named-service | nfsd-service | packet-triggered-subscribers |peer-selection-
service |pgm | pic-services-logging | pki-service |ppp | ppp-service | pppoe | protected-system-
domain-service | redundancy-interface-process | remote-operations | root-system-domain-service |
routing <logical-system logical-system-name>  | sampling | sbc-configuration-process | sdk-
service |service-deployment | services | snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd |usb-control| vrrp |web-
management>
<gracefully | immediately | soft>
```

Syntax (ACX Series Routers)

```
restart
<adaptive-services |audit-process | auto-configuration | autoinstallation |chassis-control |
class-of-service |clksyncd-service |database-replication| dhcp-service | diameter-service | disk-
monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management | ethernet-link-fault-
management |event-processing | firewall | general-authentication-service | gracefully |
immediately |interface-control | ipsec-key-management | l2-learning | lacp |link-management |mib-
process | mountd-service |mpls-traceroute |mspd | named-service | nfsd-service | pgm | pki-
service | ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
sdk-service |secure-neighbor-discovery | service-deployment | services | snmp |soft | statistics-
service| subscriber-management | subscriber-management-helper | tunnel-oamd | vrrp>
```

Syntax (EX Series Switches)

```
restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp | dhcp-
service | diameter-service | dot1x-protocol | ethernet-link-fault-management | ethernet-
switching | event-processing | firewall | general-authentication-service | interface-control |
kernel-health-monitoring | kernel-replication | l2-learning | lacp | license-service | link-
management | lldpd-service | mib-process | mountd-service | multicast-snooping | pgm |
```

```
redundancy-interface-process | remote-operations | routing | secure-neighbor-discovery | service-
deployment | sflow-service | snmp | vrrp | web-management>
```

Syntax (MX Series Routers)

```
restart
<adaptive-services | ancpd-service | application-identification | audit-process | auto-
configuration | bbe-stats-service | captive-portal-content-delivery | ce-l2tp-service | chassis-
control | class-of-service | clksyncd-service | database-replication | datapath-trace-service |
dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging |
ethernet-connectivity-fault-management | ethernet-link-fault-management | event-processing |
firewall | general-authentication-service | gracefully | iccp-service | idp-policy | immediately
|interface-control | ipsec-key-management |kernel-health-monitoring |  kernel-replication | l2-
learning | l2cpd-service | l2tp-service | l2tp-universal-edge | lacp | license-service | link-
management | local-policy-decision-function | mac-validation | mib-process | mountd-service |
mpls-traceroute | mspd | multicast-snooping |named-service | nfsd-service | packet-triggered-
subscribers |peer-selection-service | pgm | pic-services-logging | pki-service | ppp | ppp-
service | pppoe | protected-system-domain-service | redundancy-interface-process | remote-
operations | root-system-domain-service | routing | routing <logical-system logical-system-
name>  | sampling | sbc-configuration-process | sdk-service | service-deployment | services |
snmp |soft |static-subscribers |statistics-service| subscriber-management | subscriber-
management-helper | tunnel-oamd | usb-control | vrrp | web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall |
general-authentication-service | igmp-host-services | interface-control | ipsec-key-management |
isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process | named-service | network-
access-service | nstrace-process | pgm | ppp | pppoe | redundancy-interface-process | remote-
operations |logical-system-name>  | routing | sampling |secure-neighbor-discovery | service-
```

```
deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>
```

Syntax (Routing Matrix)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | disk-monitoring |
dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control |
ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | link-management
| mib-process | pgm | pic-services-logging | ppp | pppoe | redundancy-interface-process | remote-
operations | routing <logical-system logical-system-name>  | sampling | service-deployment |
snmp>
<all | all-lcc | lcc number>
<gracefully | immediately | soft>
```

Syntax (SRX Series)

```
restart
<application-identification |application-security |audit-process |commitd-service |chassis-
control | class-of-service |database-replication |datapath-trace-service |ddns |dhcp |dhcp-
service |dynamic-flow-capture |disk-monitoring |event-processing | ethernet-connectivity-fault-
management |ethernet-link-fault-management |extensible-subscriber-services |fipsd |firewall |
firewall-authentication-service |general-authentication-service |gracefully |gprs-process |idp-
policy |immediately |interface-control | ipmi |ipsec-key-management |jflow-service |jnu-
management |jnx-wmicd-service |jsrp-service |kernel-replication |l2-learning |l2cpd-service |
lacp |license-service |logical-system-service |mib-process |mountd-service |named-service |
network-security |network-security-trace |nfsd-service |ntpd-service |pgm |pic-services-logging |
profilerd |pki-service |remote-operations |rest-api |routing |sampling |sampling-route-record |
scc-chassisd |secure-neighbor-discovery |security-intelligence |security-log |services |service-
deployment |simple-mail-client-service |soft |snmp |static-routed |statistics-service |
subscriber-management |subscriber-management-helper |system-log-vital |tunnel-oamd |uac-service |
user-ad-authentication |vrrp |web-management >
```

Syntax (TX Matrix Routers)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service |
diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing
| firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-
service | lacp | link-management | mib-process |pgm | pic-services-logging | ppp | pppoe |
redundancy-interface-process | remote-operations | routing <logical-system logical-system-name>
| sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | lcc number | scc>
<gracefully | immediately | soft>
```

Syntax (TX Matrix Plus Routers)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service |
diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing
| firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-
service | lacp | link-management | mib-process | pgm | pic-services-logging | ppp | pppoe |
redundancy-interface-process | remote-operations | routing <logical-system logical-system-name>
| sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | all-sfc | lcc number | sfc number>
<gracefully | immediately | soft>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall |
general-authentication-service | igmp-host-services | interface-control | ipsec-key-management |
isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process | named-service | network-
access-service | nstrace-process | pgm | ppp | pppoe | redundancy-interface-process | remote-
operations |logical-system-name>  | routing | sampling |secure-neighbor-discovery | service-
```

```
deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>
```

Syntax (Junos OS Evolved)

```
restart (BdL2Token | aft-sysinfo | agentd | alarmd | arpd | audit-process | bcmd_evo | bfdd |
bios-manager | charonctl | chassis-control | class-of-service | clksynced | cmevod | command-
handler | command-relay | configd | ddosd | dfwd-junos-relay | diskmgmt | distributor | dot1x-
protocol | dot1xd-agent | edo | emfca | ethernet-connectivity-fault-management | ethernet-link-
fault-management | event-processing | evo-aftmand-zx | evo-cda-zx | evo-cda-zx-diag | evo-jet-
sdk-broker | evoaft-jvisiond |  fabricHub | fabspoked-fchip | fabspoked-pfe | fabtokend | fibd |
fibd-proxy | firewall | fpa | fwstatsd | gcd | hwddual | hwdfpc | hwdspmb | icmpd | idmd-dest-
usage-class | idmd-src-usage-class | idmdbd | idmdcounter | idmdfabtoken | idmdfilter |
idmdfilterterm | idmdfwgretunnel | idmdifd | idmdifl | idmdnh | idmdoffchip32 | idmdoffchip64 |
idmdonchip | dmdpolicer | idmdrtb | idmdsensor |idmdsgid | idmdstp | ifstatsd | imgd | interface-
control | jdhcpd | jinsightd | jsd | jstatsd | kfirewall-agent | l2agent | l2ald | l2cpd | l2cpd-
agent | lacp | license-check | lldpd | mem-mgmt | mfilterd | mgd | mgd-api | mgd-pfe | mgmt-ethd
| mib-process | mplsoamd | mstr | mstrzk | msvcsd | mstrzk | msvcsd | mustd | na-grpcd | na-
mqttd | ndp | netdefaultsd | nlsd | objmon | objping-server | ofp | ofp-command | opticmand |
orchestratord | packetio-zx | pccd | pci-agent | pdevmand | pfestatsd | picd | ppman | ppmd |
ppmdagent | resild | routing | rpcserviced | rpdfw | securityd | sflowd | sinetd | smartd-agent-
monitor | snmp | snmpd-subagent | svcsd | syscmdd | sysepochman | sysman | sysman-ui | trace-
relay | trace-writer | xmlproxyd | ztp)
<gracefully | immediately | soft>
```

Description

Restart a Junos OS process.

CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

For Junos OS Evolved, the `restart` command also triggers a restart of the dependent applications (apps). In order to inform you which dependent apps are being restarted the following message will be logged when the `restart` command is used:

App restarting <app name>. Related apps that may be impacted - <related-app name> . For example: Jan 14 11:42:08 RE0 sysman[5100]: SYSTEM_APP_RESTARTING_WITH_RELAPPS_EVENT: App restarting re0-ifmand. Related apps that may be impacted - aggd

Starting in Junos OS Evolved Release 20.1R1, if you specify restart *app-name* and the application is not supposed to run on the platform, the error message is as follows:

```
user@device> restart fabspoked-pfe
Restart failed for fabspoked-pfe on node re0. Application is not running.
```

The restart command expands all applications names including applications that are not required for the current platform. Therefore, a user could try to do a restart for an application that is not running for the current platform. This error message communicates that the restart failed because the application was not running on the system.

Options

none Same as gracefully.

adaptive-services (Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis (TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc (TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

all-members (MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc (TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service (Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification	(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.
application-security	(Optional) Restart the application security process.
audit-process	(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.
auto-configuration	(Optional) Restart the Interface Auto-Configuration process.
autoinstallation	(EX Series switches only) (Optional) Restart the autoinstallation process.
bbe-stats-service	(MX Series routers only) (Optional) Restart bbe-statsd, the BBE statistics collection and management process.
captive-portal-content-delivery	(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.
ce-l2tp-service	(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.
chassis-control	(Optional) Restart the chassis management process.
class-of-service	(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch’s CoS configuration.
clksyncd-service	(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).
commitd-service	(Optional) Restart the committed services.
database-replication	(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.
datapath-trace-service	(Optional) Restart the packet path tracing process.
dhcp	(EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service	(Optional) Restart the Dynamic Host Configuration Protocol process.
dialer-services	(EX Series switches only) (Optional) Restart the ISDN dial-out process.
diameter-service	(Optional) Restart the diameter process.
disk-monitoring	(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.
dlsw	(QFX Series only) (Optional) Restart the data link switching (DLSw) service.
dot1x-protocol	(EX Series switches only) (Optional) Restart the port-based network access control process.
dynamic-flow-capture	(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.
ecc-error-logging	(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.
ethernet-connectivity-fault-management	(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
ethernet-link-fault-management	(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
ethernet-switching	(EX Series switches only) (Optional) Restart the Ethernet switching process.
event-processing	(Optional) Restart the event process (eventd).
extensible-subscriber-services	(Optional) Restart the extensible subscriber services process.
fibre-channel	(QFX Series only) (Optional) Restart the Fibre Channel process.
fipsd	(Optional) Restart the fipsd services.
firewall	(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.
general-authentication-service	(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gprs-process	(Optional) Restart the General Packet Radio Service (GPRS) process.
gracefully	(Optional) Restart the software process.
iccp-service	(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.
idp-policy	(Optional) Restart the intrusion detection and prevention (IDP) protocol process.
immediately	(Optional) Immediately restart the software process.
interface-control	(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.
ipmi	(Optional) Restart the intelligent platform management interface process.
ipsec-key-management	(Optional) Restart the IPsec key management process.
isdn-signaling	(QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.
jflow-service	(Optional) Restart jflow service process.
jnu-management	(Optional) Restart jnu management process.
jnx-wmicd-service	(Optional) Restart jnx wmicd service process.
jsrp-service	(Optional) Restart the Juniper Services Redundancy Protocol (jsrdp) process, which controls chassis clustering.
kernel-health-monitoring	(Optional) Restart the Routing Engine kernel health monitoring process, which enables health parameter data to be sent from kernel components to data collection applications. When you change the polling interval through `sysctl kern.jkhmd_polling_time_secs`, you must restart the kernel health monitoring process for the new polling interval to take effect.
kernel-replication	(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.
l2-learning	(Optional) Restart the Layer 2 address flooding and learning process.
l2cpd-service	(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service (M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge (MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lacp (Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number* (TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service (EX Series switches only) (Optional) Restart the feature license management process.

link-management (TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service (EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local	(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.
local-policy-decision-function	(Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.
logical-system-service	(Optional) Restart the logical system service process.
mac-validation	(Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.
member *member-id*	(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.
mib-process	(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.
mobile-ip	(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.
mountd-service	(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.
mpls-traceroute	(Optional) Restart the MPLS Periodic Traceroute process.
mspd	(Optional) Restart the Multiservice process.
multicast-snooping	(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.
named-service	(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.
network-access-service	(QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.
network-security	(Optional) Restart the network security process.

network-security-trace	(Optional) Restart the network security trace process.
nfsd-service	(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.
ntpd-service	(Optional) Restart the Network Time Protocol (NTP) process.
packet-triggered-subscribers	(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.
peer-selection-service	(Optional) Restart the Peer Selection Service process.
pgcp-service	(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the `services pgcp gateway` option.
pgm	(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.
pic-services-logging	(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.
pki-service	(Optional) Restart the PKI Service process.
ppp	(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.
ppp-service	(Optional) Restart the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.
pppoe	(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.
profilerd	(Optional) Restart the profiler process.
protected-system-domain-service	(Optional) Restart the Protected System Domain (PSD) process.
redundancy-interface-process	(Optional) Restart the ASP redundancy process.

remote-operations	(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.
rest-api	(Optional) Restart the rest api process.
root-system-domain-service	(Optional) Restart the Root System Domain (RSD) service.
routing	(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.
routing <logical-system *logical-system-name*>	(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.
sampling	(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.
sampling-route-record	(Optional) Restart the sampling route record process.
sbc-configuration-process	(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).
scc	(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).
scc-chassisd	(Optional) Restart the scc chassisd process.
sdk-service	(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.
secure-neighbor-discovery	(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.
security-intelligence	(Optional) Restart security intelligence process.
security-log	(Optional) Restart the security log process.
sfc *number*	(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

service-deployment	(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.
services	(Optional) Restart a service.
services pgcp gateway *gateway-name*	(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the `pgcp-service` option.
sflow-service	(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.
simple-mail-client-service	(Optional) Restart the simple mail client service process.
snmp	(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.
soft	(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.
static-routed	(Optional) Restart the static routed process.
static-subscribers	(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.
statistics-service	(Optional) Restart the process that manages the Packet Forwarding Engine statistics.
subscriber-management	(Optional) Restart the Subscriber Management process.
subscriber-management-helper	(Optional) Restart the Subscriber Management Helper process.
system-log-vital	(Optional) Restart system log vital process.
tunnel-oamd	(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

uac-service	(Optional) Restart the Unified Access Control (UAC) process.
usb-control	(MX Series routers) (Optional) Restart the USB control process.
user-ad-authentication	(Optional) Restart User ad Authentication process
vrrp	(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.
web-management	(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level

reset

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```

restart interface-control gracefully

```
user@host> restart interface-control gracefully
Interface control process started, pid 41129
```

restart interface-control (Junos OS Evolved)

```
user@host> restart interface-control
interface-control restart requested
Restarted aggd on re0
Restarted ifmand on re0
```

Release Information

Command introduced before Junos OS Release 7.4.

Options added:

- `dynamic-flow-capture` in Junos OS Release 7.4.
- `dlsw` in Junos OS Release 7.5.
- `event-processing` in Junos OS Release 7.5.
- `ppp` in Junos OS Release 7.5.
- `l2ald` in Junos OS Release 8.0.
- `link-management` in Junos Release 8.0.
- `pgcp-service` in Junos OS Release 8.4.
- `sbc-configuration-process` in Junos OS Release 9.5.
- `services pgcp gateway` in Junos OS Release 9.6.
- `sfc` and `all-sfc` for the TX Matrix Router in Junos OS Release 9.6.
- Command introduced before Junos OS Release 9.2 on SRX Series devices.
- `bbe-stats-service` in Junos OS Release 18.4R1 on MX Series routers.

- `kernel-health-monitoring` in Junos OS Release 19.1R1.
- Introduced in Junos OS Evolved Release 19.1R1.

RELATED DOCUMENTATION

Overview of Operational Mode Commands

request modem wireless create-profile

IN THIS SECTION

- Syntax | 802
- Description | 802
- Options | 803
- Required Privilege Level | 803
- Sample Output | 804
- Release Information | 804

Syntax

```
request modem wireless create-profile interface-name access-point-name access-point-name
authentication-method authentication-method  profile-id profile-id sip-password sip-password sip-
user-id sip-id slot sim-slot-number
```

Description

Create a profile. The Subscriber Identity Module (SIM) uses a profile to establish a connection with the network. You can configure up to 16 profiles for each SIM card. The LTE Mini-PIM supports two SIM cards and so you can configure a total of 32 profiles, although only one profile can be active at a time.

To create a profile, you must obtain the following information from the service provider:

- Username and password
- Access point name (APN)
- Authentication (Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP))

Options

- *interface-name*—The LTE interface is `cl-x/0/0`, where *x* is the slot number in which the LTE Mini-PIM is installed.
- `access-point-name` *access-point-name*—Access point name (APN). Obtain the APN from the service provider. You can specify only a single APN in a profile.
- `authentication-method`—The authentication protocol that the SIM card uses to authenticate with the wireless network. Obtain the authentication information from the service provider. The authentication protocol used by the SIM card must match the protocol used by the service provider. The `authentication-method` can be one of the following:
  - CHAP
  - PAP
  - None
- `profile-id` *profile-id*—Profile identification number for the profile. The default value is 1. The range of possible values is from 1 through 16.
- `sip-password` *sip-password*—Simple IP password. Obtain the password from the service provider.
- `sip-user-id` *sip-id*—Simple IP user identification. Obtain the username from the service provider.
- `slot` *sim-slot-number*—The slot in which the SIM card is inserted. The value can be either 1 or 2.

Required Privilege Level

maintenance

Sample Output

request modem wireless create-profile

```
user@host> request modem wireless create-profile cl-1/0/0 access-point-name apn authentication-
method pap profile-id 2 sip-password 123 sip-user-id userid slot 1
Issued create profile request successfully.
Please use 'show modem wireless profiles' to check profile status
```

Release Information

Command introduced in Junos OS 9.5. The `slot` *sim-slot-number* option is introduced in Junos OS 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless profiles | 1068

request modem wireless fota

IN THIS SECTION

- Syntax | 805
- Description | 805
- Required Privilege Level | 805
- Sample Output | 805
- Release Information | 806

Syntax

```
request modem wireless fota interface-name (enable | disable)
```

Description

Enable or disable over-the-air (OTA) firmware upgrade for the modem on the LTE Mini-PIM. OTA firmware upgrade enables automatic and timely upgrade of modem firmware when new firmware versions are available. The OTA upgrade can be enabled or disabled on the LTE Mini-PIM. OTA is disabled by default.

Required Privilege Level

maintenance

Sample Output

request modem wireless fota (enable)

```
user@host> request modem wireless fota cl-1/0/0 enable
Set FOTA on modem succeeded
```

request modem wireless fota (disable)

```
user@host> request modem wireless fota cl-1/0/0 disable
Set FOTA on modem succeeded
```

Release Information

Command introduced in Junos OS 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless firmware | 1059

request modem wireless sim-lock

IN THIS SECTION

- Syntax | 806
- Description | 806
- Options | 807
- Required Privilege Level | 807
- Sample Output | 807
- Release Information | 807

Syntax

```
request modem wireless sim-lock enable interface-name pin pin
```

Description

Lock the Subscriber Identity Module (SIM) on the Mini-PIM. The SIM lock does not take effect until the next reboot of the services gateway. You can verify the locked mode using the `show modem wireless firmware` command.

NOTE: If there are two SIMs installed on the LTE Mini-PIM, then only the active SIM is locked. After the SIM is locked, it cannot connect to the network. The SIM must be unlocked before it is used to connect to the network.

Options

- *interface-name*—The LTE Mini-PIM is denoted as cl-*x*/0/0, where *x* is the slot number in which the LTE Mini-PIM is installed.
- `pin` *pin*—Four-digit personal identification number (PIN). Obtain the PIN from the service provider.

  NOTE: If the PIN is entered incorrectly three consecutive times, the SIM card is blocked. Obtain a PIN unblocking key (PUK) from the service provider.

Required Privilege Level

maintenance

Sample Output

request modem wireless sim-lock

```
user@host> request modem wireless sim-lock enable cl-1/0/0 pin 4321
Issued SIM 2 lock state request successfully.
Please use 'show modem wireless firmware' to check SIM status
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

request modem wireless sim-unlock | 808

request modem wireless sim-unlock

IN THIS SECTION

- Syntax | 808
- Description | 808
- Options | 809
- Required Privilege Level | 809
- Sample Output | 809
- Release Information | 809

Syntax

```
request modem wireless sim-unlock interface-name pin unlock-code
```

Description

Unlock the Subscriber Identity Module (SIM) on the LTE Mini-PIM. Some service providers lock the SIM to prevent unauthorized access to the service provider's network. If this is the case, you will need to unlock the SIM by using an personal identification number (PIN), which is provided by the service provider. You can verify the unlocked mode using the `show modem wireless firmware` command.

NOTE: If there are two SIM cards installed on the Mini-PIM, then only the active SIM card is unlocked.

The SIM must be unlocked before it can be used to connect to the service provider's network.

Options

- *interface-name*—The LTE interface is denoted as cl-*x*/0/0, where *x* is the slot number in which the LTE Mini-PIM is installed.
- `pin` *unlock-code*—Four-digit personal identification number (PIN). Obtain the PIN from the service provider.

  NOTE: If the PIN is entered incorrectly three consecutive times, the SIM card is blocked. Obtain a PIN unblocking key (PUK) from the service provider.

Required Privilege Level

maintenance

Sample Output

request modem wireless sim-unlock

```
user@host> request modem wireless sim-unlock cl-1/0/0 pin 1234
Issued SIM 2 unlock request successfully.
Please use 'show modem wireless firmware' to check SIM status
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

request modem wireless sim-lock | 806

request wlan access-point packet capture

IN THIS SECTION

- Syntax | 810
- Description | 810
- Options | 811
- Required Privilege Level | 811
- What's Next | 811
- Sample Output | 811
- Release Information | 812

Syntax

```
request wlan access-point packet-capture start ap-name interface interface-name duration
duration filename capture-file sizefile-size-max promiscuous filter-mac mac-address
```

Description

Capture the packets for the interfaces on wireless LAN interface Mini-PIM and to check whether the packets are received from Wi-Fi interface and sent to JUNOS. In chassis cluster mode, packet capture command is supported only on active wireless LAN interface.

Options

- `start` *`access-point name`*—Access point name to start the packet capture.
  - `duration`—Specify the capture duration. Range: 10 through 3600 seconds
  - `filename`—Specify the captured filename.
  - `filter-mac`—Specify the MAC address of the interface.
  - `interface`—Specify the name of the interface.
  - `promiscuous`—Enable the promiscuous mode.
  - `size`—Specify the maximum file size. Range: 64 through 4096 kilobytes.
- `stop` *`access-point name`*—Stop the packet capture.

Required Privilege Level

view

WHAT'S NEXT

Sample Output

request wlan access-point packet capture

```
user@host> request wlan access-point packet-capture start filename wlan.pcap e09-22-ha-ap
interface Radio1VAP0 duration 10
Starting packet capture

Capture interface:   Radio1VAP0

File             :   wlan.pcap

Duration         :   10 seconds
```

```
File size max    :   1024 kilobytes
```

Release Information

Command introduced in Junos OS release 20.3R1.

RELATED DOCUMENTATION

- Wi-Fi Mini-Physical Interface Module Overview | 443
- Configure Wi-Fi Mini-PIM | 446
- wlan | 767

show chassis fpc (View)

IN THIS SECTION

- Syntax | 813
- Description | 813
- Options | 813
- Required Privilege Level | 813
- Output Fields | 814
- Sample Output | 815
- Sample Output | 817
- Sample Output | 817
- Sample Output | 821
- Release Information | 822

Syntax

```
          show chassis fpc
<detail <          fpc-slot        >| <node (          node-id         | local | primary)>> |
<node (          node-id         | local | primary)> |
<pic-status <          fpc-slot        >| <node (          node-id         | local | primary)>>
```

Description

Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options

- none—Display status information for all FPCs.
- `detail`—(Optional) Display detailed FPC status information.
- *fpc-slot* —(Optional) Display information about the FPC in this slot.
- `node`—(Optional) For chassis cluster configurations, display status information for all FPCs or for the specified FPC on a specific node (device) in the cluster.
  - *node-id* —Identification number of the node. It can be 0 or 1.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.
- `pic-status`—(Optional) Display status information for all FPCs or for the FPC in the specified slot (see *fpc-slot*).

Required Privilege Level

view

Output Fields

Table 50 on page 814 lists the output fields for the `show chassis fpc` command. Output fields are listed in the approximate order in which they appear.

Table 50: show chassis fpc Output Fields

Field Name	Field Description
`Slot or Slot State`	Slot number and state. The state can be one of the following conditions: • `Dead`—Held in reset because of errors. • `Diag`—Slot is being ignored while the device is running diagnostics. • `Dormant`—Held in reset. • `Empty`—No FPC is present. • `Online`—FPC is online and running. • `Present`—FPC is detected by the device, but is either not supported by the current version of Junos OS or inserted in the wrong slot. The output also states either `Hardware Not Supported` or `Hardware Not In Right Slot`. FPC is coming up but not yet online. • `Probed`—Probe is complete; awaiting restart of the Packet Forwarding Engine (PFE). • `Probe-wait`—Waiting to be probed.
`Temp (C) or Temperature`	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.
`Total CPU Utilization (%)`	Total percentage of CPU being used by the FPC's processor.
`Interrupt CPU Utilization (%)`	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.
`Memory DRAM (MB)`	Total DRAM, in megabytes, available to the FPC's processor.

Table 50: show chassis fpc Output Fields *(Continued)*

Field Name	Field Description
Heap Utilization (%)	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.
Start Time	Time when the Routing Engine detected that the FPC was running.
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.
PIC type	(pic-status output only) Type of FPC.

Sample Output

show chassis fpc

```
user@host> show chassis fpc
                                Temp      CPU Utilization (%)     Memory    Utilization (%)
     Slot State         (C)       Total  Interrupt        DRAM (MB) Heap  Buffer
      0  Online          -------------------- CPU less FPC --------------------
      1  Online          --------------------  Not Usable --------------------
      2  Online          -------------------- CPU less FPC --------------------
```

show chassis fpc (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc
                     Temp      CPU Utilization (%)       Memory    Utilization (%)
Slot State           (C)      Total  Interrupt           DRAM (MB) Heap     Buffer
  0  Empty
```

```
  1  Empty
  2  Empty
  3  Online            37          3          0             1024        7     42
  4  Empty
  5  Empty
  6  Online            30           8         0               1024        23    30
  7  Empty
  8  Empty
  9  Empty
 10  Empty
 11  Empty
```

show chassis fpc (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc
                     Temp  CPU Utilization (%)   CPU Utilization (%)
Memory
                           Utilization (%)
Slot State            (C)  Total  Interrupt      1min   5min   15min  DRAM
(MB)
                       Heap     Buffer
  0  Online            36     20          0       20     19     19
1024
                             4         26
  1  Online            35      8          0        8      8      8
2048
                            12         14
  2  Online            40     21          0       20     20     20
3584
                             5         13
```

Sample Output

show chassis fpc detail 2

```
user@host> show chassis fpc detail 2
Slot 2 information:
  State                                   Online
  Temperature                              37
  Total CPU DRAM                         1024 MB
  Total RLDRAM                            0 MB
  Total DDR DRAM                          0 MB
  Start time:                               2012-07-18 07:18:50 PDT
  Uptime:                                   4 days, 21 hours, 51 minutes, 59 seconds
  Max Power Consumption                   0 Watts
```

Sample Output

show chassis fpc pic-status (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
Slot 3   Online       SRX5k SPC
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
Slot 6   Online       SRX5k DPC 4x 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SPC2)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
```

```
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 2   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 3   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices with SRX5K-MPC)

```
user@host> show chassis fpc pic-status

Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 6   Offline      SRX5k SPC II
Slot 9   Online       SRX5k SPC II
  PIC 0  Online       SPU Flow
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
```

```
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Online       SRX5k IOC II
  PIC 0  Online       1x 100GE CFP
  PIC 2  Online       2x 40GE QSFP+
```

show chassis fpc pic-status (SRX5600 and SRX5800 devices when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status

Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

show chassis fpc pic-status (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis fpc pic-status

node0:
--------------------------------------------------------------------------
Slot 0   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 1   Offline      SRX5k SPC II
```

```
Slot 2   Online       SRX5k DPC 4X 10GE
  PIC 0  Online       1x 10GE(LAN/WAN) RichQ
  PIC 1  Online       1x 10GE(LAN/WAN) RichQ
  PIC 2  Online       1x 10GE(LAN/WAN) RichQ
  PIC 3  Online       1x 10GE(LAN/WAN) RichQ
Slot 9   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP
  PIC 1  Online       10x 1GE(LAN) SFP
  PIC 2  Online       10x 1GE(LAN) SFP
  PIC 3  Online       10x 1GE(LAN) SFP
Slot 10  Online       SRX5k IOC II
  PIC 0  Online       10x 10GE SFP+
  PIC 2  Online       1x 100GE CFP
Slot 11  Offline      SRX5k IOC II
```

show chassis fpc pic-status (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3 and when Express Path [formerly known as services offloading] is configured)

```
user@host> show chassis fpc pic-status
Slot 0   Offline      SRX5k DPC 40x 1GE
Slot 1   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Online       SPU Flow
  PIC 3  Online       SPU Flow
Slot 2   Offline      SRX5k SPC
Slot 4   Online       SRX5k IOC3 24XGE+6XLG
  PIC 2  Online       3x 40GE QSFP+- np-cache/services-offload
  PIC 3  Online       3x 40GE QSFP+- np-cache/services-offload
Slot 5   Online       SRX5k IOC II
  PIC 0  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 1  Online       10x 1GE(LAN) SFP- np-cache/services-offload
  PIC 2  Online       10x 10GE SFP+- np-cache/services-offload
```

Sample Output

show chassis fpc pic-status for HA (SRX5600 and SRX5800 devices)

```
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow

node1:
--------------------------------------------------------------
Slot 4   Online       SRX5k DPC 40x 1GE
  PIC 0  Online       10x 1GE RichQ
  PIC 1  Online       10x 1GE RichQ
  PIC 2  Online       10x 1GE RichQ
  PIC 3  Online       10x 1GE RichQ
Slot 5   Online       SRX5k SPC
  PIC 0  Online       SPU Cp-Flow
  PIC 1  Online       SPU Flow
```

show chassis fpc pic-status for HA (SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G (IOC3) or SRX5K-MPC3-40G10G (IOC3)

```
user@host> show chassis fpc pic-status
user@host> show chassis fpc pic-status
node0:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
```

```
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline

node1:
--------------------------------------------------------------------------
Slot 2   Online       SRX5k IOC3 24XGE+6XLG
  PIC 0  Online       12x 10GE SFP+
  PIC 1  Online       12x 10GE SFP+
  PIC 2  Offline      3x 40GE QSFP+
  PIC 3  Offline      3x 40GE QSFP+
Slot 4   Online       SRX5k IOC II
  PIC 2  Online       10x 10GE SFP+
Slot 5   Online       SRX5k SPC II
  PIC 0  Online       SPU Cp
  PIC 1  Online       SPU Flow
  PIC 2  Offline
  PIC 3  Offline
```

Release Information

Command modified in Junos OS Release 9.2.

Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced.

NOTE: On SRX5K-MPC3-40G10G (IOC3), all four PICs cannot be powered on. A maximum of two PICs can be powered on at the same time. By default, PIC0 and PIC1 are online.

Use the `set chassis fpc <slot> pic <pic> power off` command to choose the PICs you want to power on.

When you use the `set chassis fpc <slot> pic <pic> power off` command to power off PIC0 and PIC1, PIC2 and PIC3 are automatically turned on.

When you switch from one set of PICs to another set of PICs using the `set chassis fpc <slot> pic <pic> power off` command again, ensure that there is 60 seconds duration between the two actions, otherwise core files are seen during the configuration.

The Table 51 on page 823 summarizes the SRX5K-MPC3-40G10G (IOC3) PICs selected for various configuration scenarios.

Table 51: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary

CLI Configuration	PIC Selection
Default (i.e. no CLI configuration)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1, PIC-2 and PIC-3 powered OFF	Online: PIC-0 Offline: PIC-1, PIC-2, PIC-3
PIC-0, PIC-2 and PIC-3 powered OFF	Online: PIC-1 Offline: PIC-0, PIC-2, PIC-3
PIC-0, PIC-1 and PIC-3 powered OFF	Online: PIC-2 Offline: PIC-0, PIC-1, PIC-3
PIC-0, PIC-1 and PIC-2 powered OFF	Online: PIC-3 Offline: PIC-0, PIC-1, PIC-2
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-2 and PIC-3 powered OFF	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3
PIC-1 and PIC-2 powered OFF	Online: PIC-0, PIC-3 Offline: PIC-1, PIC-2
PIC-0 and PIC-3 powered OFF	Online: PIC-2, PIC-1 Offline: PIC-0, PIC-3

Table 51: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary *(Continued)*

CLI Configuration	PIC Selection
PIC-0 and PIC-1 powered OFF	Online: PIC-2, PIC-3 Offline: PIC-0, PIC-1
All other combinations of PICs being powered OFF (Invalid)	Online: PIC-0, PIC-1 Offline: PIC-2, PIC-3 Default PICs will be selected for the invalid combinations. Also, a system log message will be displayed to indicate the invalid combination PIC selection.

RELATED DOCUMENTATION

Understanding Interfaces | 2

show chassis hardware (View)

IN THIS SECTION

- Syntax | 825
- Description | 825
- Options | 826
- Required Privilege Level | 826
- Output Fields | 826
- show chassis hardware | 832
- show chassis hardware (SRX4200) | 842
- show chassis hardware (vSRX 3.0) | 843
- show chassis hardware clei-models | 843
- Release Information | 844

Syntax

```
show chassis hardware
<clei-models | detail  / extensive |  models | node (node-id         | all | local | primary)>
```

Description

Display chassis hardware information.

Starting in Junos OS Release 20.1R1, when vSRX 3.0 performs resource management, the vCPUs and RAM available to the instance are assigned based on what has been allocated prior to launching the instance. A maximum of 32 cores will be assigned to SRXPFE, for flow processing. Any allocation of cores in excess of 32 will automatically be assigned to the Routing Engine. For example, if 36 cores are allocated to the VM during the creation process, 32 cores are assigned for flow processing and 4 cores will be assigned to the RE. For memory allocations, up to 64G of vRAM would be used by the SRXPFE. Any allocated memory in excess of 64G would be assigned to system memory and would not be used for maintaining flow sessions information.

Table 52: Recommended vCPU and vRAM Combinations

vCPU Number	vRAM Size (G)
2	4
5	8
9	16
17	32

On a deployed vSRX, only memory scale up is supported. Scaling down memory on a deployed vSRX, is not supported. If you need to scale down memory, then a fresh install is required.

Options

- `clei-models`—(Optional) Display Common Language Equipment Identifier Code (CLEI) barcode and model number for orderable field-replaceable units (FRUs).
- `detail` | `extensive`—(Optional) Display the specified level of output.
- `models`—(Optional) Display model numbers and part numbers for orderable FRUs.
- `node`—(Optional) For chassis cluster configurations, display chassis hardware information on a specific node (device) in the cluster.
  - *`node-id`* —Identification number of the node. It can be 0 or 1.
  - `local`—Display information about the local node.
  - `primary`—Display information about the primary node.

Required Privilege Level

view

Output Fields

Table 53 on page 826 lists the output fields for the `show chassis hardware` command. Output fields are listed in the approximate order in which they appear.

Table 53: show chassis hardware Output Fields

Field Name	Field Description
`Item`	Chassis component—Information about the backplane; power supplies; fan trays; Routing Engine; each Physical Interface Module (PIM)—reported as FPC and PIC—and each fan, blower, and impeller.
`Version`	Revision level of the chassis component.

Table 53: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
Part Number	Part number for the chassis component.
Serial Number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the device chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the device chassis.
Assb ID or Assembly ID	Identification number that describes the FRU hardware.
FRU model number	Model number of FRU hardware component.
CLEI code	Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.
EEPROM Version	ID EEPROM version used by hardware component: 0x01 (version 1) or 0x02 (version 2).

Table 53: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
`Description`	Brief description of the hardware item: • Type of power supply. • Switch Control Board (SCB) Starting with Junos OS Release 12.1X47-D15 and Junos OS Release 17.3R1, the SRX5K-SCBE (SCB2) is introduced. • There are three SCB slots in SRX5800 devices. The third slot can be used for an SCB or an FPC. When an SRX5K-SCB was used , the third SCB slot was used as an FPC. SCB redundancy is provided in chassis cluster mode. • With an SCB2, a third SCB is supported. If a third SCB is plugged in, it provides intra-chassis fabric redundancy. • The Ethernet switch in the SCB2 provides the Ethernet connectivity among all the FPCs and the Routing Engine. The Routing Engine uses this connectivity to distribute forwarding and routing tables to the FPCs. The FPCs use this connectivity to send exception packets to the Routing Engine. • Fabric connects all FPCs in the data plane. The Fabric Manager executes on the Routing Engine and controls the fabric system in the chassis. Packet Forwarding Engines on the FPC and fabric planes on the SCB are connected through HSL2 channels. • SCB2 supports HSL2 with both 3.11 Gbps and 6.22 Gbps (SerDes) link speed and various HSL2 modes. When an FPC is brought online, the link speed and HSL2 mode are determined by the type of FPC. Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-SCB3 (SCB3) with enhanced midplane is introduced. • All existing SCB software that is supported by SCB2 is supported on SCB3. • SRX5K-RE-1800X4 mixed Routing Engine use is not supported.

Table 53: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• SCB3 works with the SRX5K-MPC (IOC2), SRX5K-MPC3-100G10G (IOC3), SRX5K-MPC3-40G10G (IOC3), and SRX5K-SPC-4-15-320 (SPC2) with current midplanes and the new enhanced midplanes. • Mixed SCB use is not supported. If an SCB2 and an SCB3 are used, the system will only power on the primary Routing Engine's SCB and will power off the other SCBs. Only the SCB in slot 0 is powered on and a system log is generated. • SCB3 supports up to 400 Gbps per slot with old midplanes and up to 500 Gbps per slot with new midplanes. • SCB3 supports fabric intra-chassis redundancy. • SCB3 supports the same chassis cluster function as the SRX5K-SCB (SCB1) and the SRX5K-SCBE (SCB2), except for in-service software upgrade (ISSU) and in-service hardware upgrade (ISHU). • SCB3 has a second external Ethernet port. • Fabric bandwidth increasing mode is not supported. Starting in Junos OS 19.3R1, SRX5K-SCB4 is supported on SRX5600 and SRX5800 devices along with SRX5K-SPC3. SRX5K-SCB4: • Interoperate with SRX5K-RE3-128G, SRX5K-RE-1800X4, IOC2, IOC3, IOC4, SPC2, and SPC3. SCB4 is compatible with all midplanes and interoperate with existing PEMs, fan trays, and front panel displays. • Does not interoperate with SCB, SCB2, and SCB3. • Supports 480-Gbps link speed per slot. • Supports 1-Gigabit Ethernet interfaces speed with SRX5K-RE-1800X4 and 1-Gigabit, 2.5-Gigabit, and 10-Gigabit Ethernet speeds with SRX5K-RE3-128G. • Support ISHU and ISSU in chassis cluster.

Table 53: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• Supports fabric bandwidth mode and redundant fabric mode on SRX5600 and SRX5800 devices. The bandwidth mode is the new default mode which is necessary to configure redundant mode in setting up the chassis cluster successfully. • Type of Flexible PIC Concentrator (FPC), Physical Interface Card (PIC), Modular Interface Cards (MICs), and PIMs. • IOCs Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-MPC3-100G10G (IOC3) and the SRX5K-MPC3-40G10G (IOC3) are introduced. • IOC3 has two types of IOC3 MPCs, which have different built-in MICs: the 24x10GE + 6x40GE MPC and the 2x100GE + 4x10GE MPC. • IOC3 supports SCB3 and SRX5000 line backplane and enhanced backplane. • IOC3 can only work with SRX5000 line SCB2 and SCB3. If an SRX5000 line SCB is detected, IOC3 is offline, an FPC misconfiguration alarm is raised, and a system log message is generated. • IOC3 interoperates with SCB2 and SCB3. • IOC3 interoperates with the SRX5K-SPC-4-15-320 (SPC2) and the SRX5K-MPC (IOC2). • The maximum power consumption for one IOC3 is 645W. An enhanced power module must be used. • The IOC3 does not support the following command to set a PIC to go offline or online: `request chassis pic fpc-slot <fpc-slot> pic-slot <pic-slot> <offline | online>` . • IOC3 supports 240 Gbps of throughput with the enhanced SRX5000 line backplane. • Chassis cluster functions the same as for the SRX5000 line IOC2.

Table 53: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• IOC3 supports intra-chassis and inter-chassis fabric redundancy mode. • IOC3 supports ISSU and ISHU in chassis cluster mode. • IOC3 supports intra-FPC and and Inter-FPC Express Path (previously known as *services offloading*) with IPv4. • NAT of IPv4 and IPv6 in normal mode and IPv4 for Express Path mode. • All four PICs on the 24x10GE + 6x40GE cannot be powered on. A maximum of two PICs can be powered on at the same time. Use the `set chassis fpc <slot> pic <pic> power off` command to choose the PICs you want to power on. Fabric bandwidth increasing mode is not supported on IOC3. • SRX Clustering Module (SCM) • Fan tray Starting in Junos OS Release 19.3R1, the SRX5K-IOC4-10G and SRX5K-IOC4-MRAT line cards are supported along with SRX5K-SPC3 on the SRX5000 series devices. SRX5K-IOC4-10G: • Interoperates with SCB3, SCB4, SRX5K-RE-1800X4, SRX5K-RE3-128G, SPC2, SPC3, IOC2,IOC3, and IOC4. • Supports 480-Gbps speed. • Supports 40X10GE Interfaces with SCB3. • 40 10-Gigabit Ethernet port provides 10-Gigabit Ethernet MACsec support. • Supports reth and aggregated interfaces on the chassis cluster. • Supports ISSU and logical system on the chassis cluster. • Does not support SCB2.

Table 53: show chassis hardware Output Fields *(Continued)*

Field Name	Field Description
	• SRX5K-IOC4-MRAT with SCB3 supports 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet Interfaces. • For hosts, the Routing Engine type. Starting with Junos OS Release 12.1X47-D15 and Junos OS Release 17.3R1, the SRX5K-RE-1800X4 Routing Engine is introduced. • The SRX5K-RE-1800X4 has an Intel Quad core Xeon processor, 16 GB of DRAM, and a 128-GB solid-state drive (SSD). The number 1800 refers to the speed of the processor (1.8 GHz). The maximum required power for this Routing Engine is 90W. **NOTE**: The SRX5K-RE-1800X4 provides significantly better performance than the previously used Routing Engine, even with a single core. Starting in Junos OS Release 19.3R1, SRX5K-RE3-128G Routing Engine is supported along with SRX5K-SPC3 on the SRX5000 series devices. SRX5K-RE3-128G: • Provides improved control plane performance and scalability. SRX5K-RE3-128G has Intel's Haswell-EP based processor with six cores. • Supports two 200G SSDs to store log files and 128-GB of memory for storing routing and forwarding tables and for other routing engines. • Interoperates with SCB3, SCB4, SRX5K-RE3-128G, SPC2, SPC3, IOC2, IOC3, and IOC4. • Does not support SCB2 and SRX5K-RE-1800X4.

show chassis hardware

show chassis hardware (SRX5800)

```
user@host> show chassis hardware
node0:
```

```
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1267B0FAGA      SRX5800
Midplane         REV 42   760-063937   ACRL3065          Enhanced SRX5800
Backpl
                   ane
FPM Board        REV 05   760-061272   CAHE4860          Front Panel Display
PDM              Rev 01   740-063049   QCS2209509D       Power Distribution
Modu
                 le
PEM 0            Rev 04   740-034724   QCS171002016      PS 4.1kW; 200-240V AC
i
               n
PEM 1            Rev 11   740-027760   QCS1825N07S       PS 4.1kW; 200-240V AC
i
               n
Routing Engine 0 REV 01   750-095568   CALK8884          SRX5k RE-2000x6
Routing Engine 1 REV 01   750-095568   CADZ9076          SRX5k RE-2000x6
CB 0             REV 26   750-031391   CALV3002          SRX5k SCB4
CB 1             REV 26   750-031391   CALV3009          SRX5k SCB4
FPC 2            REV 28   750-073435   CALS4630          SPC3
  CPU                     BUILTIN      BUILTIN           SRX5k vCPP Broadwell
  PIC 0                   BUILTIN      BUILTIN           SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 17   750-044175   CABE7777          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 4            REV 08   750-061262   CAFD8147          SRX5k IOC II
  CPU            REV 02   711-061263   CAFV7488          SRX5k MPC PMB
  MIC 0          REV 03   750-055732   CAFV9369          20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 02   740-011613   PNB1GJR           SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
FPC 5            REV 10   750-062242   CAKX2328          SRX5k IOC3 2CGE+4XGE
  PIC 0                   BUILTIN      BUILTIN           2x 10GE SFP+
    Xcvr 0       REV 01   740-021308   ANA07RE           SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AQF0RBJ           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 100GE CFP2
  PIC 2                   BUILTIN      BUILTIN           2x 10GE SFP+
```

```
    Xcvr 0       REV 01   740-031980   AA1650304RF       SFP+-10G-SR
    Xcvr 1       REV 01   740-021308   AQ93BDK           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           1x 100GE CFP2
FPC 8            REV 46   750-056519   CALC4514          SRX5k IOC4 MRAT
  CPU            REV 21   750-057177   CALC3494          SMPC PMB
  PIC 0                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 0       REV 01   740-059437   000T20128         QSFP28-LPBK
    Xcvr 1       REV 01   740-067443   1ACP13450KH       QSFP+-40G-SR4
  PIC 1                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 0       REV 01   740-059437   0000T3443         QSFP28-LPBK
Fan Tray 0       REV 06   740-035409   ACAE9390          Enhanced Fan Tray
Fan Tray 1       REV 06   740-035409   ACAE9386          Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1267B01AGA      SRX5800
Midplane         REV 42   760-063937   ACRL3068          Enhanced SRX5800 Backplane
FPM Board        REV 05   760-061272   CAJX9988          Front Panel Display
PDM              Rev 01   740-063049   QCS2209507A       Power Distribution Module
PEM 0            Rev 11   740-027760   QCS1822N0EY       PS 4.1kW; 200-240V AC in
PEM 1            Rev 03   740-034724   QCS17020203F      PS 4.1kW; 200-240V AC in
Routing Engine 0 REV 01   750-095568   CALK8904          SRX5k RE-2000x6
Routing Engine 1 REV 01   750-095568   CADZ9076          SRX5k RE-2000x6
CB 0             REV 26   750-031391   CALV3010          SRX5k SCB4
CB 1             REV 26   750-031391   CALV3000          SRX5k SCB4
FPC 2            REV 28   750-073435   CAKZ9620          SPC3
  CPU                     BUILTIN      BUILTIN           SRX5k vCPP Broadwell
  PIC 0                   BUILTIN      BUILTIN           SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
FPC 3            REV 18   750-054877   CACH4082          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 4            REV 08   750-061262   CAFD8165          SRX5k IOC II
  CPU            REV 02   711-061263   CAFV7507          SRX5k MPC PMB
  MIC 0          REV 03   750-055732   CAFV6603          20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 0     REV 01   740-011613   AM0805S8M4N       SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
```

```
FPC 5            REV 03   750-062242   CAFZ2748          SRX5k IOC3 2CGE+4XGE
  PIC 0                   BUILTIN      BUILTIN           2x 10GE SFP+
    Xcvr 0       REV 01   740-021308   11T511100788      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AS92WJ0           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 100GE CFP2
  PIC 2                   BUILTIN      BUILTIN           2x 10GE SFP+
    Xcvr 0       REV 01   740-031980   AA1650304EZ       SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   ANS0EAR           SFP+-10G-SR
  PIC 3                   BUILTIN      BUILTIN           1x 100GE CFP2
FPC 8            REV 46   750-056519   CALC4526          SRX5k IOC4 MRAT
  CPU            REV 21   750-057177   CALF5727          SMPC PMB
  PIC 0                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
    Xcvr 1       REV 01   740-067443   1ACP13450L9       QSFP+-40G-SR4
  PIC 1                   BUILTIN      BUILTIN           MRATE-6xQSFPP-XGE-XLGE-CGE
Fan Tray 0       REV 06   740-035409   ACAE9298          Enhanced Fan Tray
Fan Tray 1       REV 06   740-035409   ACAE9314          Enhanced Fan Tray

{primary:node0}
```

show chassis hardware (SRX5600 and SRX5800 devices for SRX5K-MPC)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN12170EAAGA      SRX 5800
Midplane         REV 01   710-041799   ACAX3849          SRX 5800 Backplane
FPM Board        REV 01   710-024632   CAAX7297          Front Panel Display
PDM              Rev 03   740-013110   QCS170250DU       Power Distribution Modu
le
PEM 0            Rev 03   740-034724   QCS17020203F      PS 4.1kW; 200-240V AC i
n
PEM 1            Rev 03   740-034724   QCS17020203C      PS 4.1kW; 200-240V AC i
n
PEM 2            Rev 04   740-034724   QCS17100200A      PS 4.1kW; 200-240V AC i
n
PEM 3            Rev 03   740-034724   QCS17080200M      PS 4.1kW; 200-240V AC i
n
Routing Engine 0 REV 11   740-023530   9012047437        SRX5k RE-13-20
CB 0             REV 09   710-024802   CAAX7202          SRX5k SCB
CB 1             REV 09   710-024802   CAAX7157          SRX5k SCB
FPC 0            REV 07   750-044175   CAAD0791          SRX5k SPC II
```

```
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Cp
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 1           REV 07   750-044175   CAAD0751         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 2           REV 28   750-020751   CAAW1817         SRX5k DPC 4X 10GE
  CPU           REV 04   710-024633   CAAZ5269         SRX5k DPC PMB
  PIC 0                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
    Xcvr 0      REV 02   740-014289   T10A00404        XFP-10G-SR
  PIC 1                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
  PIC 2                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
  PIC 3                 BUILTIN      BUILTIN          1x 10GE(LAN/WAN) RichQ
FPC 6           REV 02   750-044175   ZY2552           SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
FPC 9           REV 10   750-044175   CAAP5932         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 10          REV 22   750-043157   ZH8192           SRX5k IOC II   CPU            REV 08
711-043360   YX3879           SRX5k MPC PMB
  MIC 0         REV 01   750-049488   YZ2084           10x 10GE SFP+
    PIC 0               BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0    REV 01   740-031980   AMB0HG3          SFP+-10G-SR
      Xcvr 1    REV 01   740-031980   AM20B6F          SFP+-10G-SR
  MIC 1         REV 19   750-049486   CAAH3504         1x 100GE CFP
    PIC 2               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01   740-035329   X000D375         CFP-100G-SR10
FPC 11          REV 07.04.07 750-043157 CAAJ8771      SRX5k IOC II   CPU            REV 08
711-043360   CAAJ3881         SRX5k MPC PMB
  MIC 0         REV 19   750-049486   CAAH0979         1x 100GE CFP
    PIC 0               BUILTIN      BUILTIN          1x 100GE CFP
      Xcvr 0    REV 01   740-035329   UP1020Z          CFP-100G-SR10
  MIC 1         REV 08   750-049487   CAAM1160         2x 40GE QSFP+
    PIC 2               BUILTIN      BUILTIN          2x 40GE QSFP+
      Xcvr 0    REV 01   740-032986   QB151094         QSFP+-40G-SR4
```

```
      Xcvr 1     REV 01   740-032986   QB160509          QSFP+-40G-SR4
Fan Tray 0       REV 04   740-035409   ACAE0875          Enhanced Fan Tray
Fan Tray 1       REV 04   740-035409   ACAE0876          Enhanced Fan Tray
```

show chassis hardware (with 20-Gigabit Ethernet MIC with SFP)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN108DA5AAGA      SRX 5800
Midplane         REV 02   710-013698   TR0037            SRX 5600 Midplane
FPM Board        REV 02   710-014974   JY4635            Front Panel Display
PDM              Rev 02   740-013110   QCS10465005       Power Distribution Module
PEM 0            Rev 03   740-023514   QCS11154040       PS 1.7kW; 200-240VAC in
PEM 2            Rev 02   740-023514   QCS10504014       PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 05   740-015113   1000681023        RE-S-1300
CB 0             REV 05   710-013385   JY4775            SRX5k SCB
FPC 1            REV 17   750-020751   WZ6349            SRX5k DPC 4X 10GE
  CPU            REV 02   710-024633   WZ0718            SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0                NON-JNPR     C724XM088         XFP-10G-SR
  PIC 1                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
    Xcvr 0       REV 02   740-011571   C831XJ08S         XFP-10G-SR
  PIC 2                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
  PIC 3                   BUILTIN      BUILTIN           1x 10GE(LAN/WAN) RichQ
FPC 3            REV 22   750-043157   ZH8189            SRX5k IOC II
  CPU            REV 06   711-043360   YX3912            SRX5k MPC PMB
  MIC 0          REV 01   750-055732   CACF9115          20x 1GE(LAN) SFP
    PIC 0                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 2     REV 02   740-013111   B358549           SFP-T
      Xcvr 9     REV 02   740-011613   PNB1FQS           SFP-SX
    PIC 1                 BUILTIN      BUILTIN           10x 1GE(LAN) SFP
      Xcvr 9     REV 02   740-011613   PNB1FFF           SFP-SX
FPC 5            REV 01   750-027945   JW9665            SRX5k FIOC
  CPU
FPC 8            REV 08   750-023996   XA7234            SRX5k SPC
  CPU            REV 02   710-024633   XA1599            SRX5k DPC PMB
  PIC 0                   BUILTIN      BUILTIN           SPU Cp-Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
```

```
Fan Tray 0       REV 03   740-014971   TP0902           Fan Tray
Fan Tray 1       REV 01   740-014971   TP0121           Fan Tray
```

show chassis hardware (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE [SCB2] and SRX5K-RE-1800X4 [RE2])

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1251EA1AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2657          Enhanced SRX5600 Midplane
FPM Board        REV 01   710-024631   CABY3551          Front Panel Display
PEM 0            Rev 03   740-034701   QCS13380901P      PS 1.4-2.6kW; 90-264V AC in
PEM 1            Rev 03   740-034701   QCS133809019      PS 1.4-2.6kW; 90-264V AC in
Routing Engine 0 REV 02   740-056658   9009210105        SRX5k RE-1800X4
Routing Engine 1 REV 02   740-056658   9013115551        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CADW3663          SRX5k SCB3
CB 1             REV 01   750-062257   CADZ3263          SRX5k SCB3
FPC 0            REV 18   750-054877   CABG6043          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 1            REV 01   750-062243   CAEE5918          SRX5k IOC3 24XGE+6XLG
  CPU            REV 02   711-062244   CADX8509          RMPC PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
    Xcvr 0       REV 01   740-031980   273363A01891      SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   273363A01915      SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   ANA0BK6           SFP+-10G-SR
    Xcvr 3       REV 01   740-031980   AP407GA           SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   MUC20G1           SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 15   750-049136   CAEE5845          MPC5E 24XGE OTN Mezz
FPC 3            REV 11   750-043157   CACL7452          SRX5k IOC II
  CPU            REV 04   711-043360   CACP1977          SRX5k MPC PMB
  MIC 0          REV 04   750-049488   CABL4759          10x 10GE SFP+
```

```
    PIC 0                 BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 0     REV 01   740-021308   CF36KM0SY        SFP+-10G-SR
      Xcvr 1     REV 01   740-021308   MUC0MF2          SFP+-10G-SR
      Xcvr 2     REV 01   740-021308   CF36KM01S        SFP+-10G-SR
      Xcvr 3     REV 01   740-021308   MUC229N          SFP+-10G-SR
FPC 5             REV 07   750-044175   CAAD0764         SRX5k SPC II
  CPU                     BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN          SPU Flow
  PIC 1                   BUILTIN      BUILTIN          SPU Flow
  PIC 2                   BUILTIN      BUILTIN          SPU Flow
  PIC 3                   BUILTIN      BUILTIN          SPU Flow
Fan Tray                                               Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FE77AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2970         Enhanced SRX5600 Midplane
FPM Board        REV 01   710-024631   CABY3552         Front Panel Display
PEM 0            Rev 03   740-034701   QCS133809028     PS 1.4-2.6kW; 90-264V AC in
PEM 1            Rev 03   740-034701   QCS133809027     PS 1.4-2.6kW; 90-264V AC in
Routing Engine 0 REV 02   740-056658   9009218294       SRX5k RE-1800X4
Routing Engine 1 REV 02   740-056658   9013104758       SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEB8180         SRX5k SCB3
CB 1             REV 01   750-062257   CADZ3334         SRX5k SCB3
FPC 0            REV 18   750-054877   CACJ9834         SRX5k SPC II
  CPU                     BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN          SPU Cp
  PIC 1                   BUILTIN      BUILTIN          SPU Flow
  PIC 2                   BUILTIN      BUILTIN          SPU Flow
  PIC 3                   BUILTIN      BUILTIN          SPU Flow
FPC 1            REV 01   750-062243   CAEB0981         SRX5k IOC3 24XGE+6XLG
  CPU            REV 02   711-062244   CAEA4644         RMPC PMB
  PIC 0                   BUILTIN      BUILTIN          12x 10GE SFP+
    Xcvr 0       REV 01   740-031980   AP41BLH          SFP+-10G-SR
    Xcvr 1       REV 01   740-031980   AQ400SL          SFP+-10G-SR
    Xcvr 2       REV 01   740-031980   AP422LJ          SFP+-10G-SR
    Xcvr 3       REV 01   740-021308   AMG0RBT          SFP+-10G-SR
    Xcvr 9       REV 01   740-021308   MUC2FRG          SFP+-10G-SR
  PIC 1                   BUILTIN      BUILTIN          12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN          3x 40GE QSFP+
  PIC 3                   BUILTIN      BUILTIN          3x 40GE QSFP+
```

```
  WAN MEZZ       REV 15   750-049136   CAEA4837          MPC5E 24XGE OTN Mezz
FPC 3            REV 11   750-043157   CACA8784          SRX5k IOC II
  CPU            REV 04   711-043360   CACA8820          SRX5k MPC PMB
  MIC 0          REV 05   750-049488   CADF0521          10x 10GE SFP+
    PIC 0                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 0     REV 01   740-030658   AD1130A00PV       SFP+-10G-USR
      Xcvr 1     REV 01   740-031980   AN40MVV           SFP+-10G-SR
      Xcvr 2     REV 01   740-021308   CF36KM37B         SFP+-10G-SR
      Xcvr 3     REV 01   740-021308   AD153830DSZ       SFP+-10G-SR
  MIC 1          REV 01   750-049487   CABB5961          2x 40GE QSFP+
    PIC 2                 BUILTIN      BUILTIN           2x 40GE QSFP+
      Xcvr 1     REV 01   740-032986   QB160513          QSFP+-40G-SR4
FPC 5            REV 02   750-044175   ZY2569            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                Enhanced Fan Tray
```

show chassis hardware (SRX5400, SRX5600, and SRX5800 devices with SRX5000 line SRX5K-SCB3 [SCB3] with enhanced midplanes and SRX5K-MPC3-100G10G [IOC3] or SRX5K-MPC3-40G10G [IOC3])

```
user@host> show chassis hardware
node0:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN1250870AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2578          Enhanced SRX5600
Midplane

FPM Board        REV 02   710-017254   KD9027            Front Panel Display
PEM 0            Rev 03   740-034701   QCS13090900T      PS 1.4-2.6kW; 90-264V
A
                     C in
PEM 1            Rev 03   740-034701   QCS13090904T      PS 1.4-2.6kW; 90-264V
A
                     C in
Routing Engine 0 REV 01   740-056658   9009196496        SRX5k RE-1800X4
```

```
CB 0             REV 01   750-062257   CAEC2501          SRX5k SCB3
FPC 0            REV 10   750-056758   CADC8067          SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Cp
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
FPC 2            REV 01   750-062243   CAEE5924          SRX5k IOC3 24XGE+6XLG
  CPU            REV 01   711-062244   CAEB4890          SRX5k IOC3 PMB
  PIC 0                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 1                   BUILTIN      BUILTIN           12x 10GE SFP+
  PIC 2                   BUILTIN      BUILTIN           3x 40GE QSFP+
    Xcvr 0       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
    Xcvr 2       REV 01   740-038623   MOC13156230449    QSFP+-40G-CU1M
  PIC 3                   BUILTIN      BUILTIN           3x 40GE QSFP+
  WAN MEZZ       REV 01   750-062682   CAEE5817          24x 10GE SFP+ Mezz
FPC 4            REV 11   750-043157   CACY1595          SRX5k IOC II
  CPU            REV 04   711-043360   CACZ8879          SRX5k MPC PMB
  MIC 1          REV 04   750-049488   CACM6062          10x 10GE SFP+
    PIC 2                 BUILTIN      BUILTIN           10x 10GE SFP+
      Xcvr 7     REV 01   740-021308   AD1439301TU       SFP+-10G-SR
      Xcvr 8     REV 01   740-021308   AD1439301SD       SFP+-10G-SR
      Xcvr 9     REV 01   740-021308   AD1439301TS       SFP+-10G-SR
FPC 5            REV 05   750-044175   ZZ1371            SRX5k SPC II
  CPU                     BUILTIN      BUILTIN           SRX5k DPC PPC
  PIC 0                   BUILTIN      BUILTIN           SPU Flow
  PIC 1                   BUILTIN      BUILTIN           SPU Flow
  PIC 2                   BUILTIN      BUILTIN           SPU Flow
  PIC 3                   BUILTIN      BUILTIN           SPU Flow
Fan Tray                                                 Enhanced Fan Tray

node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                JN124FEC0AGB      SRX5600
Midplane         REV 01   760-063936   ACRE2946          Enhanced SRX5600 Midplane
FPM Board        test     710-017254   test              Front Panel Display
PEM 0            Rev 01   740-038514   QCS114111003      DC 2.6kW Power Entry Module
PEM 1            Rev 01   740-038514   QCS12031100J      DC 2.6kW Power Entry Module
Routing Engine 0 REV 01   740-056658   9009186342        SRX5k RE-1800X4
CB 0             REV 01   750-062257   CAEB8178          SRX5k SCB3
FPC 0            REV 07   750-044175   CAAD0769          SRX5k SPC II
```

```
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Cp
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
FPC 4           REV 11   750-043157   CACY1592         SRX5k IOC II
  CPU           REV 04   711-043360   CACZ8831         SRX5k MPC PMB
  MIC 1         REV 04   750-049488   CACN0239         10x 10GE SFP+
    PIC 2               BUILTIN      BUILTIN          10x 10GE SFP+
      Xcvr 7    REV 01   740-031980   ARN23HW          SFP+-10G-SR
      Xcvr 8    REV 01   740-031980   ARN2FVW          SFP+-10G-SR
      Xcvr 9    REV 01   740-031980   ARN2YVM          SFP+-10G-SR
FPC 5           REV 10   750-056758   CADA8736         SRX5k SPC II
  CPU                   BUILTIN      BUILTIN          SRX5k DPC PPC
  PIC 0                 BUILTIN      BUILTIN          SPU Flow
  PIC 1                 BUILTIN      BUILTIN          SPU Flow
  PIC 2                 BUILTIN      BUILTIN          SPU Flow
  PIC 3                 BUILTIN      BUILTIN          SPU Flow
Fan Tray                                             Enhanced Fan Tray
```

show chassis hardware (SRX4200)

command-name

```
user@host> show chassis hardware

Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DK2816AR0020      SRX4200
Mainboard        REV 01   650-071675   16061032317       SRX4200
Routing Engine 0          BUILTIN      BUILTIN           SRX Routing Engine
FPC 0                     BUILTIN      BUILTIN           FEB
  PIC 0                   BUILTIN      BUILTIN           8x10G-SFP
    Xcvr 0       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 1       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 2       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 3       REV 01   740-038153   MOC11511530020    SFP+-10G-CU3M
    Xcvr 4       REV 01   740-021308   04DZ06A00364      SFP+-10G-SR
    Xcvr 5       REV 01   740-031980   233363A03066      SFP+-10G-SR
```

```
    Xcvr 6       REV 01   740-021308   AL70SWE           SFP+-10G-SR
    Xcvr 7       REV 01   740-031980   ALN0N6C           SFP+-10G-SR
    Xcvr 8       REV 01   740-030076   APF16220018NK1    SFP+-10G-CU1M
Power Supply 0   REV 04   740-041741   1GA26241849       JPSU-650W-AC-AFO
Power Supply 1   REV 04   740-041741   1GA26241846       JPSU-650W-AC-AFO
Fan Tray 0                                              SRX4200 0, Front to Back Airflow - AFO
Fan Tray 1                                              SRX4200 1, Front to Back Airflow - AFO
Fan Tray 2                                              SRX4200 2, Front to Back Airflow - AFO
Fan Tray 3                                              SRX4200 3, Front to Back Airflow - AFO
```

show chassis hardware (vSRX 3.0)

command-name

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                806dddb1a141      VSRX
Midplane
System IO
Routing Engine                                          VSRX-2CPU-8G memory
FPC 0                                                   FPC
  PIC 0                                                 VSRX DPDK GE
Power Supply 0
```

show chassis hardware clei-models

show chassis hardware clei-models (SRX5600 and SRX5800 devices with SRX5000 line SRX5K-SCBE [SCB2] and SRX5K-RE-1800X4 [RE2])

```
user@host> show chassis hardware clei-models node 1
node1:
--------------------------------------------------------------------------
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
```

```
Midplane          REV 01   710-024803                SRX5800-BP-A
FPM Board         REV 01   710-024632                SRX5800-CRAFT-A
PEM 0             Rev 04   740-034724                SRX5800-PWR-4100-AC
PEM 1             Rev 05   740-034724                SRX5800-PWR-4100-AC
Routing Engine 0  REV 01   740-056658   COUCATTBAA   SRX5K-RE-1800X4
CB 0              REV 01   750-056587   COUCATSBAA   SRX5K-SCBE
CB 1              REV 01   750-056587   COUCATSBAA   SRX5K-SCBE
CB 2              REV 01   750-056587   COUCATSBAA   SRX5K-SCBE
FPC 0             REV 18   750-054877   COUCATLBAA   SRX5K-SPC-4-15-320
  CPU                      BUILTIN
FPC 1             REV 18   750-054877   COUCATLBAA   SRX5K-SPC-4-15-320
  CPU                      BUILTIN
FPC 2             REV 18   750-054877   COUCATLBAA   SRX5K-SPC-4-15-320
  CPU                      BUILTIN
FPC 3             REV 11   750-043157   COUIBCWBAA   SRX5K-MPC
  MIC 0           REV 05   750-049486   COUIBCYBAA   SRX-MIC-1X100G-CFP
  MIC 1           REV 04   750-049488   COUIBCXBAA   SRX-MIC-10XG-SFPP
FPC 4             REV 18   750-054877   COUCATLBAA   SRX5K-SPC-4-15-320
  CPU                      BUILTIN
FPC 7             REV 18   750-054877   COUCATLBAA   SRX5K-SPC-4-15-320
  CPU                      BUILTIN
FPC 8             REV 11   750-043157   COUIBCWBAA   SRX5K-MPC
  MIC 0           REV 05   750-049486   COUIBCYBAA   SRX-MIC-1X100G-CFP
FPC 9             REV 18   750-054877   COUCATLBAA   SRX5K-SPC-4-15-320
  CPU                      BUILTIN
FPC 10            REV 18   750-054877   COUCATLBAA   SRX5K-SPC-4-15-320
  CPU                      BUILTIN
Fan Tray 0        REV 04   740-035409                SRX5800-HC-FAN
Fan Tray 1        REV 04   740-035409                SRX5800-HC-FAN
```

Release Information

Command introduced in Junos OS Release 9.2. Command modified in Junos OS Release 9.2 to include `node` option.

RELATED DOCUMENTATION

- Understanding Traffic Processing on Security Devices
- Interface Naming Conventions

show ethernet-switching mac-learning-log

IN THIS SECTION

- Syntax | 845
- Description | 845
- Required Privilege Level | 845
- Output Fields | 845
- Sample Output | 848
- Release Information | 851

Syntax

```
show ethernet-switching mac-learning-log
```

Description

Displays the event log of learned MAC addresses.

Required Privilege Level

view

Output Fields

Output fields for EX Series switches:

The following table lists the output fields for the `show ethernet-switching mac-learning-log` command. Output fields are listed in the approximate order in which they appear.

Table 54: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs.
MAC	Learned MAC address.
Deleted | Added	MAC address deleted or added to the MAC learning log.
Blocking	The forwarding state of the interface: • **blocked**—Traffic is not being forwarded on the interface. • **unblocked**—Traffic is forwarded on the interface.
`Flags`	Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.

Output fields for QFX Series switches, QFabric, NFX Series devices and EX4600:

Table 55 on page 846 lists the output fields for the `show ethernet-switching mac-learning-log` command. Output fields are listed in the approximate order in which they appear.

Table 55: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp in UTC when the MAC operation occurred.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs. The name of the VLAN on which the MAC is learned.
MAC	Learned MAC address.

Table 55: show ethernet-switching mac-learning-log Output Fields *(Continued)*

Field Name	Field Description
`Event op`	MAC address that are added, learned, deleted, changed or moved from one interface to another interface.
`Interface Name`	The name of the interface on which the MAC address is learned. When a MAC address is moved, there is another field with the name of the interface. The log displays the name of the interface from where the MAC address moved, and the name of the interface to where the MAC address moved.
`Flags`	Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.

Output fields for SRX Series devices:

Table 56 on page 847 lists the output fields for the show ethernet-switching mac-learning-log command on SRX Series devices. Output fields are listed in the approximate order in which they appear.

Table 56: show ethernet-switching-mac-learning-log Output Fields

Field Name	Field Description
`Date and Time`	Timestamp when the MAC address was added or deleted from the log.
`VLAN-IDX`	VLAN index. An internal value assigned by Junos OS for each VLAN.
`MAC`	Learned MAC address.
`Deleted | Added`	MAC address deleted or added to the MAC learning log.
`Blocking`	The forwarding state of the interface: • blocked—Traffic is not being forwarded on the interface. • unblocked—Traffic is forwarded on the interface.

Sample Output

show ethernet-switching mac-learning-log (EX Series switch)

```
user@switch> show ethernet-switching mac-learning-log
Mon Feb 25 08:07:05 2008
 vlan_name v1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v9 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v3 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v12 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v13 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v3 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:05:00:00:05 was learned
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:30:48:90:54:89 was learned
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:5e:00:01:00 was learned
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:5e:00:01:08 was learned
[output truncated]
```

show ethernet-switching mac-learning-log (QFX Series Switches, QFabric, NFX Series Devices and EX4600)

```
user@switch> show ethernet-switching mac-learning-log
Mon Jun 30 13:49:49 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on ge-1/0/22.0 with
flags: 0x2001f     << MAC address that as dynamically learned
Mon Jun 30 13:50:29 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was deleted from ge-1/0/22.0
with flags: 0x1080  << MAC address that was deleted
Mon Jun 30 13:51:28 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was added to ge-1/0/22.0 with
flags: 0x2013f        << Static MAC address that was added
Mon Jun 30 13:51:46 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was deleted from ge-1/0/22.0
with flags: 0x1120   << delete of Static MAC address that was deleted
Mon Jun 30 13:52:03 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on ge-1/0/22.0 with
flags: 0x2001f    << MAC address that was dynamically learned
Mon Jun 30 13:52:11 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was moved from ge-1/0/22.0 to
ge-1/0/21.0 with flags: 0x2101f    << MAC address that was moved
Mon Jun 30 13:54:24 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was changed on ge-1/0/21.0 with
flags: 0x2113f   << MAC address that changed from a dynamic address to a static address
```

show ethernet-switching mac-learning-log (SRX Series devices)

```
user@host> show ethernet-switching mac-learning-log
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was deleted
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was deleted
```

```
Wed Mar 18 08:07:05 2009
vlan_idx 4 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 9 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 10 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 11 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 13 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 14 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 15 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:00 was added
Wed Mar 18 08:07:05 2009
vlan_idx 18 mac 00:00:5E:00:53:AA was learned
Wed Mar 18 08:07:05 2009
vlan_idx 5 mac 00:00:5E:00:53:AB was learned
Wed Mar 18 08:07:05 2009
vlan_idx 6 mac 00:00:5E:00:53:AC was learned
Wed Mar 18 08:07:05 2009
vlan_idx 16 mac 00:00:5E:00:53:AD was learned
Wed Mar 18 08:07:05 2009
vlan_idx 7 mac 00:00:5E:00:53:AE was learned
Wed Mar 18 08:07:05 2009
vlan_idx 8 mac 00:00:5E:00:53:AF was learned
Wed Mar 18 08:07:05 2009
vlan_idx 12 mac 00:00:5E:00:53:AG was learned
[output truncated]
```

Release Information

Command introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

- *show ethernet-switching table*
- *show ethernet-switching interfaces*
- *show ethernet-switching table*
- *show ethernet-switching interfaces*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- Example: Connecting an EX Series Access Switch to a Distribution Switch

show ethernet-switching table

IN THIS SECTION

- Syntax (QFX Series, QFabric, NFX Series and EX4600) | **852**
- Syntax (EX Series) | **852**
- Syntax (EX Series, MX Series and QFX Series) | **852**
- Syntax (SRX Series) | **853**
- Description | **853**
- Options | **853**
- Additional Information | **855**
- Required Privilege Level | **855**
- Output Fields | **855**
- Sample Output | **860**
- Sample Output | **880**
- Sample Output | **881**
- Sample Output | **882**

Sample Output | 883
Release Information | 884

Syntax (QFX Series, QFabric, NFX Series and EX4600)

```
show ethernet-switching table
<brief | detail | extensive | summary>
<interface interface-name>
<management-vlan>
<sort-by (name | tag)>
<vlan vlan-name>
```

Syntax (EX Series)

```
show ethernet-switching table
<brief | detail | extensive | summary>
<interface interface-name>
<management-vlan>
<persistent-mac <interface interface-name>>
<sort-by (name | tag)>
<vlan vlan-name>
```

Syntax (EX Series, MX Series and QFX Series)

```
show ethernet-switching table
<brief | count | detail | extensive | summary>
<address>
<instance instance-name>
<interface interface-name>
isid isid
<logical-system logical-system-name>
```

```
<persistent-learning (interface interface-name | mac mac-address)>
<address>
<vlan-id (all-vlan | vlan-id)>
<vlan-name (all | vlan-name)>
```

Syntax (SRX Series)

```
show ethernet-switching table (brief |detail |extensive) interface  interface-name
```

Description

Displays the Ethernet switching table.

(MX Series routers,EX Series switches only) Displays Layer 2 MAC address information.

Options

For QFX Series, QFabric, NFX Series and EX4600:

none	(Optional) Display brief information about the Ethernet switching table.
brief | detail | extensive | summary	(Optional) Display the specified level of output.
interface *interface-name*	(Optional) Display the Ethernet switching table for a specific interface.
management-vlan	(Optional) Display the Ethernet switching table for a management VLAN.
persistent-mac <interface *interface-name*>	(Optional) Display the persistent MAC addresses learned for all interfaces or a specified interface. You can use this command to view entries that you want to clear for an interface that you intentionally disabled.
sort-by (*name* **|** *tag***)**	(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

vlan *vlan-name*	(Optional) Display the Ethernet switching table for a specific VLAN.

For EX Series, MX Series and QFX Series:

none	Display all learned Layer 2 MAC address information.
brief | count | detail | extensive | summary	(Optional) Display the specified level of output.
address	(Optional) Display the specified learned Layer 2 MAC address information.
instance *instance-name*	(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.
interface *interface-name*	(Optional) Display learned Layer 2 MAC addresses for the specified interface.
isid *isid*	(Optional) Display learned Layer 2 MAC addresses for the specified ISID.
logical-system *logical-system-name*	(Optional) Display Ethernet-switching statistics information for the specified logical system.
persistent-learning (interface *interface-name* **| mac** *mac-address***)**	(Optional) Display dynamically learned MAC addresses that are retained despite device restarts and interface failures for a specified interface, or information about a specified MAC address.
vlan-id (all-vlan | *vlan-id***)**	(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.
vlan-name (all | *vlan-name***)**	(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.

For SRX Series:

- `none`—(Optional) Display brief information about the Ethernet switching table.
- `brief` | `detail` | `extensive`—(Optional) Display the specified level of output.
- `interface-name`—(Optional) Display the Ethernet switching table for a specific interface.

Additional Information

When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.

Required Privilege Level

view

Output Fields

For QFX Series, QFabric, NFX Series and EX4600:

The following table lists the output fields for the `show ethernet-switching table` command on QFX Series, QFabric, NFX Series and EX4600. Output fields are listed in the approximate order in which they appear.

Table 57: show ethernet-switching table Output Fields on QFX, NFX Series and EX4600

Field Name	Field Description	Level of Output
VLAN	Name of a VLAN.	All levels
MAC address	MAC address associated with the VLAN.	All levels
Type	Type of MAC address: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members.	All levels
Age	Time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels

Table 57: show ethernet-switching table Output Fields on QFX, NFX Series and EX4600 *(Continued)*

Field Name	Field Description	Level of Output
Interfaces	Interface associated with learned MAC addresses or with the **All-members** option (flood entry).	All levels
Learned	For learned entries, the time at which the entry was added to the Ethernet switching table.	**detail**, **extensive**

For EX Series switches:

The following table lists the output fields for the `show ethernet-switching table` command on EX Series switches. Output fields are listed in the approximate order in which they appear.

Table 58: show ethernet-switching table Output Fields on EX Series Switches

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels
Tag	The VLAN ID tag name or number.	**extensive**
MAC or **MAC address**	The MAC address associated with the VLAN.	All levels
Type	The type of MAC address. Values are: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members. • **persistent**—The learned MAC addresses that will persist across restarts of the switch or interface-down events.	All levels except **persistent-mac**

Table 58: show ethernet-switching table Output Fields on EX Series Switches *(Continued)*

Field Name	Field Description	Level of Output
Type	The type of MAC address. Values are: • **installed**—addresses that are in the Ethernet switching table. • **uninstalled**—addresses that could not be installed in the table or were uninstalled in an interface-down event and will be reinstalled in the table when the interface comes back up.	**persistent-mac**
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** (flood entry).	All levels
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.	**detail**, **extensive**
Nexthop index	The next-hop index number.	**detail**, **extensive**
persistent-mac	**installed** indicates MAC addresses that are in the Ethernet switching table and **uninstalled** indicates MAC addresses that could not be installed in the table or were uninstalled in an interface-down event (and will be reinstalled in the table when the interface comes back up).	

For EX Series, MX Series and QFX Series:

The table describes the output fields for the `show ethernet-switching table` command on EX Series, MX Series and QFX Series. Output fields are listed in the approximate order in which they appear.

Table 59: show ethernet-switching table Output fields on EX, MX and QFX Series

Field Name	Field Description
Routing instance	Name of the routing instance.
VLAN name	Name of the VLAN.

Table 59: show ethernet-switching table Output fields on EX, MX and QFX Series *(Continued)*

Field Name	Field Description
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: • **S**—Static MAC address is configured. • **D**—Dynamic MAC address is configured. • **L**—Locally learned MAC address is configured. • **SE**—MAC accounting is enabled. • **NM**—Non-configured MAC. • **R**—Locally learned MAC address is configured. • **O**—OVSDB learned MAC address is configured.
Age	This field is not supported.
Logical interface	Name of the logical interface. Name of the VTEP logical interface learned over remote VTEP.
SVLBNH/VENH Index	**NOTE**: This field appears on QFX5*XXX* switches that support dynamic load balancing in an EVPN-VXLAN network. Next-hop index number associated with the MAC address of a multihomed remote device in an EVPN-VXLAN network. This index number appears when the Logical Interface column displays esi.*nnnn*. The index number can be an SVLBNH, a VENH, or a remote virtual tunnel endpoint (VTEP). To get more information about SVLBNHs, VENHs, and remote VTEPs, see *show ethernet-switching vxlan-tunnel-end-point svlbnh*.
Active source	IP address or Ethernet segment identifier (ESI) of remote entity on which MAC address is learned.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.

Table 59: show ethernet-switching table Output fields on EX, MX and QFX Series *(Continued)*

Field Name	Field Description
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or VLAN in which the MAC address was learned.
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning-tree-protocolepoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

For SRX Series:

Table 60 on page 859 lists the output fields for the `show ethernet-switching table` command. Output fields are listed in the approximate order in which they appear.

Table 60: show ethernet-switching table Output Fields on SRX Series

Field Name	Field Description
`VLAN`	The name of a VLAN.
`MAC address`	The MAC address associated with the VLAN.

Table 60: show ethernet-switching table Output Fields on SRX Series *(Continued)*

Field Name	Field Description
Type	The type of MAC address. Values are: • static—The MAC address is manually created. • learn—The MAC address is learned dynamically from a packet's source MAC address. • flood—The MAC address is unknown and flooded to all members.
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
Interfaces	Interface associated with learned MAC addresses or All-members (flood entry).
Learned	For learned entries, the time which the entry was added to the Ethernet switching table.

Sample Output

show ethernet-switching table (Enhanced Layer 2 Software on QFX Series, QFabric, NFX Series and EX460)

```
user@switch> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan1               b0:c6:9a:ca:3c:01   D             -    ae1.0
    vlan1               b0:c6:9a:ca:3c:03   D             -    ae1.0
```

```
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan10              b0:c6:9a:ca:3c:01   D             -    ae1.0
    vlan10              b0:c6:9a:ca:3c:03   D             -    ae1.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan2               b0:c6:9a:ca:3c:01   D             -    ae1.0
    vlan2               b0:c6:9a:ca:3c:03   D             -    ae1.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    vlan3               b0:c6:9a:ca:3c:01   D             -    ae1.0
    vlan3               b0:c6:9a:ca:3c:03   D             -    ae1.0

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)


Ethernet switching table : 2 entries, 2 learned
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
```

```
    vlan4               b0:c6:9a:ca:3c:01   D             -   ae1.0
    vlan4               b0:c6:9a:ca:3c:03   D             -   ae1.0
```

show ethernet-switching table (QFX Series switches, SVLBNH/VENH field)

```
user@switch> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 16 entries, 16 learned Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                SVLBNH/      Active
   name                address             flags    interface              VENH Index   source
   vlanBLACK           00:00:5e:00:53:01   DR       esi.1773               1782
05:00:00:02:9a:00:00:00:68:00
   vlanBLACK           00:00:5e:00:53:02   DR       esi.1773               1782
05:00:00:02:9a:00:00:00:68:00
   vlanBLACK           00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
   vlanBLACK           00:00:5e:00:53:00   D        vtep.32769
10.0.1.1
   vlanBLUE            00:00:5e:00:53:01   DR       esi.1772               1782
05:00:00:02:9a:00:00:00:66:00
   vlanBLUE            00:00:5e:00:53:02   DR       esi.1772               1782
05:00:00:02:9a:00:00:00:66:00
   vlanBLUE            00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
   vlanBLUE            00:00:5e:00:53:00   D        vtep.32769
10.0.1.1
   vlanGREEN           00:00:5e:00:53:01   DR       esi.1774               1782
05:00:00:02:9a:00:00:00:67:00
   vlanGREEN           00:00:5e:00:53:02   DR       esi.1774               1782
05:00:00:02:9a:00:00:00:67:00
   vlanGREEN           00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
   vlanGREEN           00:00:5e:00:53:00   D        vtep.32769
10.0.1.1
   vlanRED             00:00:5e:00:53:01   DR       esi.1771               1782
05:00:00:02:9a:00:00:00:65:00
   vlanRED             00:00:5e:00:53:02   DR       esi.1771               1782
05:00:00:02:9a:00:00:00:65:00
```

```
   vlanRED             00:00:5e:00:53:80   D        vtep.32772
10.0.0.1
```

show ethernet-switching table (QFX Series, QFabric, NFX Series and EX460)(QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, and QFX10016)

```
user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 xe-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 xe-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 xe-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 xe-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 xe-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
[output truncated]
```

show ethernet-switching table (Private VLANs on QFX Series, QFabric, NFX Series and EX460)

```
user@switch> show ethernet-switching table
Ethernet-switching table: 10 entries, 3 learned
  VLAN              MAC address       Type         Age Interfaces
  pvlan             *                 Flood          - All-members
  pvlan             00:10:94:00:00:02 Replicated       - xe-0/0/28.0
  pvlan             00:10:94:00:00:35 Replicated       - xe-0/0/46.0
  pvlan             00:10:94:00:00:46 Replicated       - xe-0/0/4.0
  c2                *                 Flood          - All-members
  c2                00:10:94:00:00:02 Learn          0 xe-0/0/28.0
  c1                *                 Flood          - All-members
  c1                00:10:94:00:00:46 Learn          0 xe-0/0/4.0
  __pvlan_pvlan_xe-0/0/46.0__ *       Flood          - All-members
  __pvlan_pvlan_xe-0/0/46.0__ 00:10:94:00:00:35 Learn        0 xe-0/0/46.0
```

show ethernet-switching table brief (QFX Series, QFabric, NFX Series and EX460)

```
user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 xe-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 xe-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 xe-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 xe-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
```

```
  T111              00:19:e2:50:ac:00 Learn          0 xe-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 xe-0/0/46.0
[output truncated]
```

show ethernet-switching table detail (QFX Series, QFabric, NFX Series and EX460)

```
user@switch> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
  F2, *
    Interface(s): xe-0/0/44.0
    Type: Flood
    Nexthop index: 0

  F2, 00:00:05:00:00:03
    Interface(s): xe-0/0/44.0
    Type: Learn, Age: 0, Learned: 2:03:09
    Nexthop index: 0

  F2, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, *
    Interface(s): xe-0/0/47.0
    Type: Flood
    Nexthop index: 0

  Linux, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
```

```
  Nexthop index: 0

Linux, 00:30:48:90:54:89
  Interface(s): xe-0/0/47.0
  Type: Learn, Age: 0, Learned: 2:03:08
  Nexthop index: 0

T1, *
  Interface(s): xe-0/0/46.0
  Type: Flood
  Nexthop index: 0

T1, 00:00:05:00:00:01
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:00:5e:00:01:00
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T1, 00:19:e2:50:63:e0
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, *
  Interface(s): xe-0/0/46.0
  Type: Flood
  Nexthop index: 0

T10, 00:00:5e:00:01:09
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, 00:19:e2:50:63:e0
```

```
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T10, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T111, *
    Interface(s): xe-0/0/15.0
    Type: Flood
    Nexthop index: 0
[output truncated]
```

show ethernet-switching table extensive (QFX Series, QFabric, NFX Series and EX460)

```
user@switch> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
  F2, *
    Interface(s): xe-0/0/44.0
    Type: Flood
    Nexthop index: 0

  F2, 00:00:05:00:00:03
    Interface(s): xe-0/0/44.0
    Type: Learn, Age: 0, Learned: 2:03:09
    Nexthop index: 0

  F2, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  Linux, *
    Interface(s): xe-0/0/47.0
    Type: Flood
    Nexthop index: 0

  Linux, 00:19:e2:50:7d:e0
    Interface(s): Router
```

```
  Type: Static
  Nexthop index: 0

Linux, 00:30:48:90:54:89
  Interface(s): xe-0/0/47.0
  Type: Learn, Age: 0, Learned: 2:03:08
  Nexthop index: 0

T1, *
  Interface(s): xe-0/0/46.0
  Type: Flood
  Nexthop index: 0

T1, 00:00:05:00:00:01
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:00:5e:00:01:00
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T1, 00:19:e2:50:63:e0
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, *
  Interface(s): xe-0/0/46.0
  Type: Flood
  Nexthop index: 0

T10, 00:00:5e:00:01:09
  Interface(s): Router
  Type: Static
  Nexthop index: 0
```

```
  T10, 00:19:e2:50:63:e0
    Interface(s): xe-0/0/46.0
    Type: Learn, Age: 0, Learned: 2:03:08
    Nexthop index: 0

  T10, 00:19:e2:50:7d:e0
    Interface(s): Router
    Type: Static
    Nexthop index: 0

  T111, *
    Interface(s): xe-0/0/15.0
    Type: Flood
    Nexthop index: 0
[output truncated]
```

show ethernet-switching table interface (QFX Series, QFabric, NFX Series and EX460)

```
user@switch> show ethernet-switching table interface xe-0/0/1
Ethernet-switching table: 1 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  V1                *                 Flood          - All-members
  V1                00:00:05:00:00:05 Learn          0 xe-0/0/1.0
```

show ethernet-switching table (EX Series switches)

```
user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 15 learned, 2 persistent
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 ge-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Persistent     0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Persistent     0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
```

```
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 ge-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table brief (EX Series switches)

```
user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 15 learned, 2 persistent entries
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 ge-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Persistent     0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Persistent     0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
```

```
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 ge-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table detail (EX Series switches)

```
user@switch> show ethernet-switching table detail
Ethernet-switching table: 5 entries, 2 learned entries
 VLAN: default, Tag: 0, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/11.0, ge-0/0/20.0, ge-0/0/30.0, ge-0/0/36.0, ge-0/0/3.0
  Type: Flood
  Nexthop index: 1307

 VLAN: default, Tag: 0, MAC: 00:1f:12:30:b8:83, Interface: ge-0/0/3.0
   Type: Learn, Age: 0, Learned: 20:09:26
   Nexthop index: 1315

 VLAN: v1, Tag: 101, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/31.0
  Type: Flood
  Nexthop index: 1313

 VLAN: v1, Tag: 101, MAC: 00:1f:12:30:b8:89, Interface: ge-0/0/31.0
  Type: Learn, Age: 0, Learned: 20:09:25
  Nexthop index: 1312
```

```
 VLAN: v2, Tag: 102, MAC: *, Interface: All-members
  Interfaces:
   ae0.0
  Type: Flood
  Nexthop index: 1317
```

show ethernet-switching table extensive (EX Series switches)

```
user@switch> show ethernet-switching table extensive
Ethernet-switching table: 3 entries, 1 learned, 5 persistent entries

    VLAN: v1, Tag: 10, MAC: *, Interface: All-members
    Interfaces:
                ge-0/0/14.0, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,
                ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/10.0,
                ge-0/0/0.0
    Type: Flood
    Nexthop index: 567

    VLAN: v1, Tag: 10, MAC: 00:21:59:c6:93:22, Interface: Router
    Type: Static
    Nexthop index: 0

    VLAN: v1, Tag: 10, MAC: 00:21:59:c9:9a:4e, Interface: ge-0/0/14.0
    Type: Learn, Age: 0, Learned: 18:40:50
    Nexthop index: 564
```

show ethernet-switching table persistent-mac (EX Series switches)

```
user@switch> show ethernet-switching table persistent-mac
VLAN              MAC address       Type       Interface
default           00:10:94:00:00:02 installed    ge-0/0/42.0
default           00:10:94:00:00:03 installed    ge-0/0/42.0
default           00:10:94:00:00:04 installed    ge-0/0/42.0
default           00:10:94:00:00:05 installed    ge-0/0/42.0
default           00:10:94:00:00:06 installed    ge-0/0/42.0
default           00:10:94:00:05:02 uninstalled ge-0/0/16.0
default           00:10:94:00:06:03 uninstalled ge-0/0/16.0
default           00:10:94:00:07:04 uninstalled ge-0/0/16.0
```

show ethernet-switching table persistent-mac interface ge-0/0/16.0 (EX Series switches)

```
VLAN              MAC address       Type       Interface
default           00:10:94:00:05:02 uninstalled ge-0/0/16.0
default           00:10:94:00:06:03 uninstalled ge-0/0/16.0
default           00:10:94:00:07:04 uninstalled ge-0/0/16.0
```

show ethernet-switching table (EX Series, MX Series and QFX Series)

```
user@host> show ethernet-switching table
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN101             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN102             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
```

```
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN104             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1101            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1102            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1103            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1104            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
```

```
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1105            00:1f:12:32:f5:c1   D             -    ae0.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1106            00:1f:12:32:f5:c1   D             -    ae0.0
[...output truncated...]
```

show ethernet-switching table brief

```
user@host> show ethernet-switching table brief
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN101             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN102             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)
```

```
Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN103             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN104             88:e0:f3:bb:07:f0   D             -    ae20.0


MAC flags (S - static MAC, D - dynamic MAC, L - locally learned
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Routing instance : default-switch
    Vlan                MAC                 MAC         Age    Logical
    name                address             flags              interface
    VLAN1101            00:1f:12:32:f5:c1   D             -    ae0.0
[...output truncated...]
```

show ethernet-switching table count

```
user@host> show ethernet-switching table count
0 MAC address learned in routing instance default-switch VLAN VLAN1000
ae26.0:1000

1 MAC address learned in routing instance default-switch VLAN VLAN101
ae20.0:101

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           101              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN102
ae20.0:102

  MAC address count per learn VLAN within routing instance:
```

```
    Learn VLAN ID      MAC count      Static MAC count
            102              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN103
ae20.0:103

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
            103              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN104
ae20.0:104

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
            104              1                 0

0 MAC address learned in routing instance default-switch VLAN VLAN105
ae20.0:105

0 MAC address learned in routing instance default-switch VLAN VLAN106
ae20.0:106

0 MAC address learned in routing instance default-switch VLAN VLAN107
ae20.0:107

0 MAC address learned in routing instance default-switch VLAN VLAN108
ae20.0:108

0 MAC address learned in routing instance default-switch VLAN VLAN109
ae20.0:109

0 MAC address learned in routing instance default-switch VLAN VLAN110
ae20.0:110

1 MAC address learned in routing instance default-switch VLAN VLAN1101
ae0.0:1101

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
           1101              1                 0

1 MAC address learned in routing instance default-switch VLAN VLAN1102
```

```
ae0.0:1102

  MAC address count per learn VLAN within routing instance:
    Learn VLAN ID      MAC count      Static MAC count
          1102              1                 0
[...output truncated...]
```

show ethernet-switching table extensive

```
user@host> show ethernet-switching table extensive

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 101
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 102
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 103
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 88:e0:f3:bb:07:f0
  Routing instance: default-switch
VLAN ID: 104
   Learning interface: ae20.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
```

```
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1101
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1102
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1103
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008

MAC address: 00:1f:12:32:f5:c1
  Routing instance: default-switch
VLAN ID: 1104
   Learning interface: ae0.0
   Layer 2 flags: in_hash,in_ifd,in_ifl,in_vlan,in_rtt,kernel,in_ifbd
   Epoch: 0                          Sequence number: 2
   Learning mask: 0x00000008
```

Sample Output

show ethernet-switching table detail (SRX Series)

```
user@host> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:5E:00:53:AC
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:00:5E:00:53:AA
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:00:5E:00:53:AB
Interface(s): Router
Type: Static
Linux, 00:00:5E:00:53:AC
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:5E:00:53:AD
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AE
Interface(s): Router
Type: Static
T1, 00:00:5E:00:53:AF
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AG
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
```

```
Type: Flood
T10, 00:00:5E:00:53:AH
Interface(s): Router
Type: Static
T10, 00:00:5E:00:53:AI
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:00:5E:00:53:AJ
Interface(s): Router
Type: Static
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table extensive (SRX Series)

```
user@host> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): ge-0/0/44.0
Type: Flood
F2, 00:00:5E:00:53:AC
Interface(s): ge-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
F2, 00:00:5E:00:53:AA
Interface(s): Router
Type: Static
Linux, *
Interface(s): ge-0/0/47.0
Type: Flood
Linux, 00:00:5E:00:53:AB
Interface(s): Router
Type: Static
Linux, 00:00:5E:00:53:AC
Interface(s): ge-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
```

```
T1, *
Interface(s): ge-0/0/46.0
Type: Flood
T1, 00:00:5E:00:53:AD
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AE
Interface(s): Router
Type: Static
T1, 00:00:5E:00:53:AF
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
T1, 00:00:5E:00:53:AG
Interface(s): Router
Type: Static
T10, *
Interface(s): ge-0/0/46.0
Type: Flood
T10, 00:00:5E:00:53:AH
Interface(s): Router
Type: Static
T10, 00:00:5E:00:53:AI
Interface(s): ge-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
T10, 00:00:5E:00:53:AJ
Interface(s): Router
Type: Static
T111, *
Interface(s): ge-0/0/15.0
Type: Flood
[output truncated]
```

Sample Output

show ethernet-switching table interface ge-0/0/1 (SRX Series)

```
user@host> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
VLAN     MAC address         Type     Age Interfaces
```

```
V1       *                   Flood    - All-members
V1       00:00:5E:00:53:AF   Learn    0 ge-0/0/1.0
```

Sample Output

show ethernet-switching table vlan-name v100 (QFX Series)

```
user@host> show ethernet-switching table vlan-name v100
MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static
           SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Pseudo mac
Ethernet switching table : 11 entries, 11 learned
Routing instance : default-switch
   Vlan                MAC                 MAC      Logical                Active
   name                address             flags    interface              source
   v100                00:00:00:00:02:80   S,NM     vtep.1074102274
112.1.1.1
   v100                00:00:00:00:03:80   S,NM     vtep.1074102275
115.1.1.1
   v100                00:00:00:00:04:80   S,NM     vtep.1074102276
116.1.1.1
   v100                00:00:00:00:05:80   S,NM     vtep.1074102277
114.1.1.1
   v100                00:00:00:00:06:80   S,NM     vtep.1074102278
101.1.1.1
   v100                00:00:00:00:07:80   S,NM     vtep.1074102279
102.1.1.1
   v100                00:00:00:00:08:80   S,NM     vtep.1074102280
103.1.1.1
   v100                00:00:00:00:09:80   S,NM     vtep.1074102281
104.1.1.1
   v100                00:00:00:00:0a:80   S,NM     vtep.1074102282
113.1.1.1
```

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 9.5 for SRX Series.

Options **summary**, **management-vlan**, and **vlan** ***vlan-name*** introduced in Junos OS Release 9.6 for EX Series switches.

Option **sort-by** and field name **tag** introduced in Junos OS Release 10.1 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Output for private VLANs introduced in Junos OS Release 12.1 for the QFX Series.

Option **persistent-mac** introduced in Junos OS Release 11.4 for EX Series switches.

Command introduced in Junos OS Release 12.3R2.

Command introduced in Junos OS Release 12.3R2 for EX Series switches.

Options `logical-system`, `persistent-learning`, and `summary` introduced in Junos OS Release 13.2X50-D10 (ELS).

Output for shared VXLAN load balancing next hop (SVLBNH) and VXLAN encapsulated next hop (VENH) introduced in Junos OS Release 20.3R1 for QFX Series switches.

Output for pseudo VTEP logical interfaces introduced in Junos OS Release 20.3R1 for QFX Series switches.

RELATED DOCUMENTATION

- Example: Setting Up Basic Bridging and a VLAN on Switches
- Example: Setting Up Bridging with Multiple VLANs
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch
- Example: Setting Up Bridging with Multiple VLANs for EX Series Switches
- Example: Setting Up Q-in-Q Tunneling on EX Series Switches
- *Dynamic Load Balancing in an EVPN-VXLAN Network*
- clear ethernet-switching table
- show ethernet-switching mac-learning-log

show igmp-snooping route (View)

IN THIS SECTION

- Syntax | 885
- Description | 885
- Options | 885
- Required Privilege Level | 886
- Output Fields | 886
- Sample Output | 886
- Release Information | 887

Syntax

```
show igmp-snooping route ( brief | detail | ethernet-switching | inet | vlan)
```

Description

Display IGMP snooping route information.

Options

- `none`—Display general parameters.
- `brief` | `detail`—(Optional) Display the specified level of output.
- `ethernet-switching`—(Optional) Display Ethernet switching information.
- `inet`—(Optional) Display inet information.
- `vlan` *vlan-id* | *vlan-name*—(Optional) Display route information for the specified VLAN.

Required Privilege Level

view

Output Fields

Table 61 on page 886 lists the output fields for the `show igmp-snooping route` command. Output fields are listed in the approximate order in which they appear.

Table 61: show igmp-snooping route Output Fields

Field Name	Field Description
`VLAN`	Name of the VLAN.
`Group`	Multicast group address.
`Next-hop`	ID associated with the next-hop device.

Sample Output

show igmp-snooping route

```
user@host> show igmp-snooping route
VLAN       Group          Next-hop
v11         203.0.113.0, *   533
Interfaces: ge-0/0/13.0, ge-0/0/1.0
v12         203.0.113.1, *   534
Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route vlan v1

```
user@host> show igmp-snooping route vlan v1
Table: 0
VLAN      Group              Next-hop
v1        203.0.113.2, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.3, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.4, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.5, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.6, *       1266
Interfaces: ge-0/0/0.0
v1        203.0.113.6, *      1266
Interfaces: ge-0/0/0.0
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Understanding Interfaces | 2

show interfaces

IN THIS SECTION

- Syntax (Gigabit Ethernet) | 888
- Syntax (10 Gigabit Ethernet) | 888

- Syntax (ACX5448, ACX5448-D, ACX710) | **889**
- Syntax (QFX5130-32CD) | **889**
- Syntax (SRX Series Devices and (vSRX and vSRX 3.0 platforms)) | **889**
- Description | **890**
- Options | **891**
- Additional Information | **892**
- Required Privilege Level | **893**
- Output Fields | **893**
- Sample output for G.fast and Annex J support | **955**
- Sample Output Gigabit Ethernet | **972**
- Sample Output | **984**
- Release Information | **1022**

Syntax (Gigabit Ethernet)

```
show interfaces ge-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (10 Gigabit Ethernet)

```
show interfaces xe-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (ACX5448, ACX5448-D, ACX710)

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (QFX5130-32CD)

```
show interfaces et-fpc/pic/port
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (SRX Series Devices and (vSRX and vSRX 3.0 platforms))

```
show interfaces (
    <interface-name>
    <brief | detail | extensive | terse>
    <controller interface-name>|
    <descriptions interface-name>|
    <destination-class (all / destination-class-name logical-interface-name)>|
    <diagnostics optics interface-name>|
    <far-end-interval interface-fpc/pic/port>|
    <filters interface-name>|
    <flow-statistics interface-name>|
    <interval interface-name>|
    <load-balancing  (detail | interface-name)>|
    <mac-database mac-address mac-address>|
    <mc-ae id identifier unit number revertive-info>|
    <media interface-name>|
```

```
    <policers interface-name>|
    <queue both-ingress-egress egress forwarding-class forwarding-class ingress 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>|
    <dsl-sfp-options (adsl-options | gfast-options | vdsl-options) >
)
```

Description

Display status information about the specified Gigabit Ethernet interface.

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Display the IPv6 interface traffic statistics about the specified Gigabit Ethernet interface for MX series routers. The input and output bytes (bps) and packets (pps) rates are not displayed for IFD and local traffic.

Display status information and statistics about interfaces on SRX Series, vSRX, and vSRX 3.0 platforms running Junos OS.

SRX4600 supports 40-Gigabit Ethernet breakouts only in PIC mode. Use the `show interfaces extensive` command to view the speed configured for the interface on SRX4600. Reboot the device for the changed configuration to take effect.

On SRX Series appliances, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

Starting in Junos OS Release 18.4R1, Output fields `Next-hop` and `vpls-status` is displayed in the `show interfaces` *`interface name`* `detail` command, only for Layer 2 protocols on MX480 routers.

In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the `show interfaces` *`interface-name`*`<media><extensive>` command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed. In

the earlier Junos OS releases, incorrect autonegotiation status was displayed even when autonegotiation was disabled.

QFX5130-32CD switches does not display the `Filters statistics` when the `show interfaces extensive` command is executed due to interface-level filter statistics related hardware limitations. See "show interfaces extensive (QFX5130-32CD)" on page 1015.

Starting in Junos OS Release 20.4R1, we support G.fast and Annex J specification with SFP xDSL for ADSL2/ADSL2+ and all VDSL2 profiles on SRX380, SRX300, SRX320, SRX340, and SRX345 devices.

Options

For Gigabit interfaces:

ge-*fpc*/*pic*/*port*	Display standard information about the specified Gigabit Ethernet interface.
	NOTE: Interfaces with different speeds are named uniformly with ge-0/0/x for backward compatibility. Use the `show interfaces` command to view the interface speeds.
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).

Additional Information

In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level

view

Output Fields

Table 62 on page 893 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 63 on page 941.

Table 62: 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*.	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

Table 62: show interfaces (Gigabit Ethernet) 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
`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: • `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*.	All levels
`Interface flags`	Information about the interface. Possible values are described in the “Interface Flags” section under *Common Output Fields Description*.	All levels

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	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 (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: *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

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`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`
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the physical interface. • `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. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the *show interfaces* command.	`detail extensive`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Errors`—Sum of the incoming frame terminated and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that 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 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminated and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. **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. • `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.	`extensive`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `Resource errors`—Sum of transmit drops.	
`Egress queues`	Total number of egress queues supported on the specified interface. **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	`detail extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `Queued packets`—Number of queued packets. • `Transmitted packets`—Number of transmitted packets. • `Dropped packets`—Number of packets dropped by the ASIC's RED mechanism. **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.	`detail extensive`
`Ingress queues`	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	`extensive`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters (Ingress)`	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • `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 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`. • `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	(On MX Series devices) Status of the `interface-transmit-statistics` configuration: Enabled or Disabled. • `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`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`OTN FEC statistics`	The forward error correction (FEC) counters provide the following statistics: • `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`	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • `Bit errors`—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • `Errored blocks`—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	`detail extensive`
`PRBS Statistics`	Displays the Pseudo Random Binary Sequence (PRBS) statistics. The PRBS Statistics are displayed in extensive, detailed, media and normal output except terse. The output is displayed per serdes lane. The output consists of the total number of iterations in error, the total number of iterations and the number of monitored seconds. An error iteration is one in which at least one bit error is seen.	`detail extensive`
`PRBS Pattern`	Specifies the pattern type, that is in the range from 7 to 58.	`detail extensive`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

Link Degrade	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.
	• Link Monitoring—Indicates if physical link degrade monitoring is enabled on the interface.
	  • 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.
	  • 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 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

MAC statistics	Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:
	• Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. 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:
	  • 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.
		extensive

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `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. **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. • `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 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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). **NOTE**: On PTX Series routers, ignore any values displayed for Input DA rejects. Input DA rejects is not supported on PTX Series routers. • `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.	`extensive`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• `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.	
`PMA PHY`	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • `Seconds`—Number of seconds the defect has been active. • `Count`—Number of times that the defect has gone from inactive to active. • `State`—State of the error. Any state other than `OK` indicates a problem. Subfields are: • `PHY Lock`—Phase-locked loop • `PHY Light`—Loss of optical signal	`extensive`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. Subfields are: • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)	extensive

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. Subfields are: • BIP-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 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information: • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. Subfields are: • BIP-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 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner status`—`OK` when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• `Link partner`—Information from the remote Ethernet device:
	  • `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:
	  • `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
		`extensive`

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	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).	
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: • Destination slot—FPC slot number.	extensive

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `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*.	All levels

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags. • 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	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
ACI VLAN	Information displayed for agent circuit identifier (ACI) interface set configured with the agent-circuit-id autoconfiguration stanza. Dynamic Profile—Name of the dynamic profile that defines the ACI interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information. **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.	brief detail extensive none
Line Identity	Information displayed for access-line-identifier (ALI) interface sets configured with the line-identity autoconfiguration stanza. • 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. If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information. **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.	detail
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under *Common Output Fields Description*.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Neighbor Discovery Protocol (NDP)Queue Statistics	NDP statistics for protocol inet6 under logical interface statistics. • 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
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

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Maximum labels`	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	`detail extensive` none
`Traffic statistics`	Number and rate of bytes and packets received and transmitted on the specified interface set. • `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`
`Transit statistics`	Number and rate of bytes and packets transiting the switch. **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.	`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

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	`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*.	`detail extensive` none
`protocol-family`	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*.	`detail extensive` none
`Destination`	IP address of the remote side of the connection.	`detail extensive` none

Table 62: show interfaces (Gigabit Ethernet) Output Fields *(Continued)*

Field Name	Field Description	Level of Output
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*.	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: • 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*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	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	Number and rate of bytes and packets received and transmitted on the physical interface. • 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`	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • `Errors`—Sum of the incoming frame terminated and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • `L3 incompletes`—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the `ignore-l3-incompletes` statement. • `L2 channel errors`—Number of times the software did not find a valid logical interface for an incoming frame. • `L2 mismatch timeouts`—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • `FIFO errors`—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • `Resource errors`—Sum of transmit drops.	`extensive`

Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame terminated and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops.	extensive

Egress queues	Total number of egress queues supported on the specified interface. **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 scenorio, 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	detail extensive
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • 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	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • 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 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. • 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	Active OTN alarms identified on the interface.	detail extensive
OTN defects	OTN defects received on the interface.	detail extensive
OTN FEC Mode	The FECmode configured on the interface. • efec—Enhanced forward error correction (EFEC) is configured to defect 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	OTN mode. • 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	The forward error correction (FEC) counters for the DWDM OTN PIC. • 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	OTN FEC excessive or degraded error alarms triggered on the interface. • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect.	detail extensive
OTN OC	OTN OC defects triggered on the interface. • 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	OTN OTU defects detected on the interface • 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`	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • `Bit errors`—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • `Errored blocks`—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	`detail` `extensive`

MAC statistics	Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following: • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • 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	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • 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.	extensive

	• 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.	
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem.	extensive
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. Subfields are: • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)	extensive

WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. Subfields are: • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)	extensive

WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. Subfields are: • BIP-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	Information about link autonegotiation.
	• Negotiation status:
	  • Incomplete—Ethernet interface has the speed or link mode configured.
	  • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• Link partner:
	  • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex.
	  • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive).
	  • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline.
	• Local resolution—Information from the link partner:
	  • 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	(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.	extensive
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • Destination slot—FPC slot number.	extensive
CoS information	Information about the CoS queue for the physical interface. • 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*.	All levels
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags. • 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:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • 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*.	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	Number and rate of bytes and packets received and transmitted on the specified interface set. • 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	Number and rate of bytes and packets transiting the switch. **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.	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*.	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*.	detail extensive none
protocol-family	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*.	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 63: 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 64 on page 942 lists the output fields for the `show interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 64: 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

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • 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
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

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`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. These fields can contain the value `None` or `Link`. • `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`	Number and rate of bytes and packets received and transmitted on the physical interface. • `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 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Input errors`	Input errors on the interface. • `Errors`—Sum of the incoming frame terminated and FCS errors. • `Drops`—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Framing errors`—Number of packets received with an invalid frame checksum (FCS). • `Runts`—Number of frames received that are smaller than the runt threshold. • `Policed discards`—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that 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 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Output errors`	Output errors on the interface. • `Carrier transitions`—Number of times the interface has gone from `down` to `up`. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • `Errors`—Sum of the outgoing frame terminated and FCS errors. • `Drops`—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • `Collisions`—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; 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`

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Queue counters and queue number`	CoS queue number and its associated user-configured forwarding class name. • `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 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output

MAC statistics	Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:
	• Total octets and total packets—Total number of octets and packets.
	• 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:
	  • 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
		extensive

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	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.”	

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Filter statistics`	`Receive` and `Transmit` statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • `Input packet count`—Number of packets received from the MAC hardware that the filter processed. • `Input packet rejects`—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • `Input DA rejects`—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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	`extensive`

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	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.	

`Autonegotiation information`	Information about link autonegotiation.
	• `Negotiation status:`
	  • `Incomplete`—Ethernet interface has the speed or link mode configured.
	  • `No autonegotiation`—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • `Complete`—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
		`extensive`

`Packet Forwarding Engine configuration`	Information about the configuration of the Packet Forwarding Engine: • `Destination slot`—FPC slot number.	`extensive`

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`CoS information`	Information about the CoS queue for the physical interface. • `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`	Status of the `interface-transmit-statistics` configuration: Enabled or Disabled.	`detail extensive`
`Queue counters (Egress)`	CoS queue number and its associated user-configured forwarding class name. • `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`
Logical Interface		
`Logical interface`	Name of the logical interface.	All levels

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`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 and transmitted on the specified interface set. • `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`	Number and rate of bytes and packets transiting the switch. **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.	`extensive`

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Security	Security zones that interface belongs to.	extensive
Flow Input statistics	Statistics on packets received by flow module.	extensive
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

Table 64: show interfaces Output Fields *(Continued)*

Field Name	Field Description	Level of Output
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 for G.fast and Annex J support

show interfaces (SRX380, SRX300, SRX320, SRX340, and SRX345)

```
user@host>  show interfaces ge-0/0/8
Physical interface: ge-0/0/8, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 520
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Link-mode: Full-duplex, Speed: 1000mbps,
BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled, Auto-negotiation: Enabled, Remote fault: Online
  DSL SFP Status:
    Chip Type             : xDSL
    Chip Firmware Version : 1_62_8463
    Training Status       : Showtime
    Training Mode         : ADSL2PLUS
    Annex Type            : Annex J
    Profile Type          : NA
    Carrier Set           : NA
    Line Status           : No Defect
  DSL SFP Statistics:                   XTU-R (DS)          XTU-C (US)
    Packet Count            :                    0                   0
    CRC Error Count         :                    0                   0
    Electrical Length (dB)  :                    0                   0
    Net Data Rate (Kbps)    :                25737                3143
```

```
  SNR Margin (dB)          :                  100                  -2
  CV Count                 :                    0                   0
  ES Count                 :                    0                   0
  SES Count                :                    0                   0
  UAS Count                :                    0                   0
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 4c:16:fc:de:30:89, Hardware address: 4c:16:fc:de:30:89
Last flapped   : 2020-10-28 19:56:29 PDT (3d 23:07 ago)
Input rate     : 20544 bps (42 pps)
Output rate    : 20544 bps (42 pps)
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
Ethernet FEC statistics              Errors
  FEC Corrected Errors                    0
  FEC Uncorrected Errors                  0
  FEC Corrected Errors Rate               0
  FEC Uncorrected Errors Rate             0
Interface transmit statistics: Disabled

Logical interface ge-0/0/8.0 (Index 77) (SNMP ifIndex 538)
  Flags: Up SNMP-Traps 0x0 VLAN-Tag [ 0x8100.10 ]  Encapsulation: ENET2
  Input packets : 0
  Output packets: 105040
  Security: Zone: trust
  Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp
  ospf ospf3 pgm pim rip ripng router-discovery rsvp sap vrrp dhcp finger ftp
  tftp ident-reset http https ike netconf ping reverse-telnet reverse-ssh
  rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
  lsping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap
  sdwan-appqoe l3-ha
  Protocol inet, MTU: 1500
  Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 0,
  Curr new hold cnt: 0, NH drop cnt: 0
    Flags: Sendbcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.1.3/24, Local: 10.1.3.2, Broadcast: 10.1.3.255
```

show interfaces (G.fast related information on SRX380, SRX300, SRX320, SRX340, and SRX345)

```
user@host>  show interfaces ge-0/0/8
Physical interface: ge-0/0/8, Enabled, Physical link is Up
G.fast mode, DS Speed: 400Mbps, US Speed: 400Mbps
Cont.....................................
.........................................
```

show interfaces terse (ACX5448, ACX5448-D, ACX710 channelized interface)

```
user@host>  show interfaces terse et-0/1/2
Interface        Admin Link Proto      Local Remote
       et-0/1/2:0             up                down
       et-0/1/2:1             up                down
       et-0/1/2:2             up                down
       et-0/1/2:3             up                down
```

show interfaces (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
```

```
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push     0x8100.512
0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Input packets : 0
    Output packets: 0
    Protocol ccc, MTU: 1522
      Flags: Is-Primary
```

show interfaces (Gigabit Ethernet on MX Series Routers)

```
user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None, Loopback:
Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault:
Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 0
  Current address: 00:00:5e:00:53:c0, Hardware address: 00:00:5e:00:53:76
  Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
        Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Input packets : 10232
    Output packets: 10294
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255    Protocol
inet6, MTU: 1500
    Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold cnt: 4, NH drop
```

```
cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 2001:db8:/32, Local: 2001:db8::5
      Addresses, Flags: Is-Preferred
        Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

show interfaces (link degrade status)

```
user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
  Interface index: 157, SNMP ifIndex: 537
  Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU Error: None, Loopback:
Disabled, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
  Last flapped   : 2014-06-18 02:36:38 PDT (02:50:50 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Link Degrade* :
  Link Monitoring                   :  Enable
  Link Degrade Set Threshold:       :  1E-7
  Link Degrade Clear Threshold:     :  1E-12
  Estimated BER                     :  1E-7
  Link-degrade event                :  Seconds     Count    State
                                           782         1     Defect Active
```

show interfaces et-0/0/0 (25-Gigabit Ethernet interfaces on PTX Series for default FEC) (Junos OS Evolved Release)

```
user@host> show interfaces et-0/0/0
Physical interface: et-0/0/0, Enabled, Physical link is Up
  Interface index: 1007, SNMP ifIndex: 503
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 25Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled, Media type: Fiber
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 84:03:28:eb:d4:44, Hardware address: 84:03:28:eb:d4:44
  Last flapped   : 2021-05-03 13:23:03 PDT (01:05:00 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC Mode  :                  FEC91  <<< Default FEC setting starting 21.1R1-EVO
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                    0
    FEC Uncorrected Errors                  0
    FEC Corrected Errors Rate               0
    FEC Uncorrected Errors Rate             0
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                   :  Disable
  Logical interface et-0/0/0.16386 (Index 1003) (SNMP ifIndex 611)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces extensive (link degrade status) (PTX10001-36MR)

```
user@host> show interfaces et-0/0/1 extensive
Physical interface: et-0/0/1, Enabled, Physical link is Down
  Interface index: 1017, SNMP ifIndex: 519, Generation: 712964572820
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Media type: Fiber
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x80
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 40:de:ad:28:7a:0a, Hardware address: 40:de:ad:28:7a:0a
  Last flapped   : 2020-08-27 12:05:18 IST (00:50:56 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :     2239638274139000         392152080440 bps
   Output bytes  :                    0                    0 bps
   Input  packets:        2239638274049             49019008 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 1, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                                0                    0                    0
    1                                0                    0                    0
    2                                0                    0                    0
    3                                0                    0                    0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK, LOCAL-DEGRADE
```

```
PCS statistics                      Seconds
  Bit errors                              0
  Errored blocks                          0
Ethernet FEC Mode  :                FEC119
Ethernet FEC statistics             Errors
  FEC Corrected Errors         166615427699
  FEC Uncorrected Errors                 12
  FEC Corrected Errors Rate         3687323
  FEC Uncorrected Errors Rate             0
MAC statistics:                     Receive         Transmit
  Total octets              2239648398609518                0
  Total packets                2239648398691                0
  Unicast packets              2239648398036                0
  Broadcast packets                        0                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total errors                             1                0
Filter statistics:
  Input packet count                       0
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot:  ( )
CoS information:
  Direction :
Interface transmit statistics: Disabled
Link Degrade :
  Link Monitoring                   :  Enable
  Link Degrade Set Threshold        :  1E-5
  Link Degrade Clear Threshold      :  1E-10
```

```
    Link Degrade War Set Threshold    :  1E-9
    Link Degrade War Clear Threshold  :  1E-11
    Estimated BER                     :  1E-5
    Link-degrade event                :  Seconds             Count               State
                                         3054                1                   Defect Active
```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```
user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 530, Generation: 154
  Interface flags: SNMP-Traps Internal: 0x4000
   Output bytes  :         240614363944            772721536 bps
   Output packets:           3538446506              1420444 pps
    Direction : Output
  Interface transmit statistics: Enabled

  Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
     Output bytes  :         195560312716            522726272 bps
     Output packets:           4251311146              1420451 pps

user@host> show interfaces ge-5/2/0.0 statistics detail
Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :              271524
     Output bytes  :            37769598
     Input  packets:                3664
     Output packets:              885790
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :            16681118
      Input  packets:                   0
      Output packets:              362633
    Local statistics:
     Input  bytes  :              271524
     Output bytes  :              308560
     Input  packets:                3664
```

```
     Output packets:                 3659
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :             37461038                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:               882131                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                    0                 0 bps
      Output bytes  :             16681118                 0 bps
      Input  packets:                    0                 0 pps
      Output packets:               362633                 0 pps
```

show interfaces brief (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

  Logical interface ge-3/0/2.0
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push     0x8100.512
0x8100.513)
    Encapsulation: VLAN-CCC
    ccc

  Logical interface ge-3/0/2.32767
    Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
```

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
```

```
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:7c, Hardware address: 00:00:5e:00:53:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Drop   bytes  :                   0                   0 bps
   Drop   packets:                   0                   0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Active alarms  : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push 0x8100.512
0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                   0
```

```
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Local statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Transit statistics:
    Input  bytes  :                   0                  0 bps
    Output bytes  :                   0                  0 bps
    Input  packets:                   0                  0 pps
    Output packets:                   0                  0 pps
  Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
    Flags: Is-Primary

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ]  Encapsulation: ENET2
  Traffic statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Local statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Transit statistics:
    Input  bytes  :                   0                  0 bps
    Output bytes  :                   0                  0 bps
    Input  packets:                   0                  0 pps
    Output packets:                   0                  0 pps
```

show interfaces extensive (Gigabit Ethernet IQ2)

```
user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70, Generation: 171
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
```

```
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags     : None
CoS queues     : 8 supported, 4 maximum usable queues
Schedulers     : 256
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:74, Hardware address: 00:00:5e:00:53:74
Last flapped   : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :          38910844056                 7952 bps
 Output bytes  :              7174605                 8464 bps
 Input  packets:            418398473                   11 pps
 Output packets:                78903                   12 pps
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Ingress traffic statistics at Packet Forwarding Engine:
 Input  bytes  :          38910799145                 7952 bps
 Input  packets:            418397956                   11 pps
 Drop   bytes  :                   0                    0 bps
 Drop   packets:                   0                    0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort            418390823            418390823                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                7133                 7133                    0
Egress queues: 4 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                 1031                 1031                    0
  1 expedited-fo                   0                    0                    0
```

```
  2 assured-forw                  0                   0                   0
  3 network-cont              77872               77872                   0
Active alarms  : None
Active defects : None
MAC statistics:                      Receive         Transmit
  Total octets                  38910844056          7174605
  Total packets                   418398473            78903
  Unicast packets               408021893366             1026
  Broadcast packets                      10               12
  Multicast packets              418398217            77865
  CRC/Align errors                       0                0
  FIFO errors                            0                0
  MAC control frames                     0                0
  MAC pause frames                       0                0
  Oversized frames                       0
  Jabber frames                          0
  Fragment frames                        0
  VLAN tagged frames                     0
  Code violations                        0  OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58
  Payload Type: 0x08
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x08
Filter statistics:
  Input packet count             418398473
  Input packet rejects                479
  Input DA rejects                    479
  Input SA rejects                      0
  Output packet count                               78903
  Output packet pad count                               0
  Output packet error count                             0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
      Remote fault: OK
  Local resolution:
      Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
```

```
    Direction : Output
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95              0      low    none
    3 network-control         5       50000000     5              0      low    none
    Direction : Input
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95              0      low    none
    3 network-control         5       50000000     5              0      low    none

  Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               812400
     Output bytes  :              1349206
     Input  packets:                 9429
     Output packets:                 9449
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :               812400
     Output bytes  :              1349206
     Input  packets:                 9429
     Output packets:                 9449
    Transit statistics:
     Input  bytes  :                   0                 7440 bps
     Output bytes  :                   0                 7888 bps
     Input  packets:                   0                   10 pps
     Output packets:                   0                   11 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 169, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
```

```
      Output Filters: F2-ge-3/0/1.0-out (53)
      Destination: 203.0.113/24, Local: 203.0.113.2, Broadcast: 203.0.113.255,
        Generation: 196
    Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the `show interfaces` command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface `Transit statistics` fields in Table 62 on page 893.

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     %           usec
    0 best-effort            95     9500000000    95              0      low    none
    3 network-control         5      500000000     5              0      low    none
    ...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```
user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional: Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```
user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
```

```
    Input  bytes  :                    0                    0 bps
    Output bytes  :      322891152287160           9627472888 bps
    Input  packets:                    0                    0 pps
    Output packets:         328809727380              1225492 pps

...

  Filter statistics:
    Output packet count           328810554250
    Output packet pad count                  0
    Output packet error count                0
...

  Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)
    Flags: SNMP-Traps Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :      322891152287160
     Input  packets:                    0
     Output packets:         328809727380
     IPv6 transit statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :      322891152287160           9627472888 bps
     Input  packets:                    0                    0 pps
     Output packets:         328809727380              1225492 pps
     IPv6 transit statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
    Protocol inet, MTU: 1500, Generation: 147, Route table: 0
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255, Generation: 141
    Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```
user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :      322857456303482           9627496104 bps
   Output bytes  :                   0                    0 bps
   Input  packets:         328775413751              1225495 pps
   Output packets:                   0                    0 pps

...

  Filter statistics:
    Input packet count            328775015056
    Input packet rejects                     1
    Input DA rejects                         0

...

  Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :      322857456303482
     Output bytes  :                   0
```

```
      Input  packets:         328775413751
      Output packets:                    0
      IPv6 transit statistics:
       Input  bytes  :                   0
       Output bytes  :                   0
       Input  packets:                   0
       Output packets:                   0
    Local statistics:
     Input  bytes  :                     0
     Output bytes  :                     0
     Input  packets:                     0
     Output packets:                     0
    Transit statistics:
     Input  bytes  :      322857456303482           9627496104 bps
     Output bytes  :                     0                    0 bps
     Input  packets:         328775413751              1225495 pps
     Output packets:                     0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                    0
      Output bytes  :                    0
      Input  packets:                    0
      Output packets:                    0
    Protocol inet, MTU: 1500, Generation: 145, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255, Generation: 139
    Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

Sample Output

Sample Output SRX Gigabit Ethernet

```
user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
```

```
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
    Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Security: Zone: public
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255
```

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: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255
```

show interfaces (Gigabit Ethernet for vSRX and vSRX 3.0)

```
user@host> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Link-mode: Half-duplex, Speed: 1000mbps,
BPDU Error: None, Loop Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:50:56:93:ef:25, Hardware address: 00:50:56:93:ef:25
  Last flapped   : 2019-03-29 01:57:45 UTC (00:00:41 ago)
  Input rate     : 1120 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
```

show interfaces detail (Gigabit Ethernet)

```
user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps, BPDU Error:
None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
```

```
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                   0                   0 bps
 Output bytes  :                   0                   0 bps
 Input  packets:                   0                   0 pps
 Output packets:                   0                   0 pps
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0 best-effort                    0                    0                    0
  1 expedited-fo                   0                    0                    0
  2 assured-forw                   0                    0                    0
  3 network-cont                   0                    0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  1                   expedited-forwarding
  2                   assured-forwarding
  3                   network-control
Active alarms  : LINK
Active defects : LINK
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
```

```
 Output packets:                    0                    0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :                    0
  ICMP packets :                    0
  VPN packets :                     0
  Multicast packets :               0
  Bytes permitted by policy :       0
  Connections established :         0
Flow Output statistics:
  Multicast packets :               0
  Bytes permitted by policy :       0
Flow error statistics (Packets dropped due to):
  Address spoofing:                 0
  Authentication failed:            0
  Incoming NAT errors:              0
  Invalid zone received packet:     0
  Multiple user authentications:    0
  Multiple incoming NAT:            0
  No parent for a gate:             0
  No one interested in self packets: 0
  No minor session:                 0
  No more sessions:                 0
  No NAT gate:                      0
  No route present:                 0
  No SA for incoming SPI:           0
  No tunnel found:                  0
  No session for a gate:            0
  No zone or NULL zone binding      0
  Policy denied:                    0
  Security association not active:  0
  TCP sequence number out of window: 0
  Syn-attack protection:            0
  User authentication errors:       0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
  Flags: Sendbcast-pkt-to-re
  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: 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       50000000     5              0      low    none
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit statistics:
   Input  bytes  :                   0                  0 bps
   Output bytes  :                   0                  0 bps
   Input  packets:                   0                  0 pps
   Output packets:                   0                  0 pps
  Security: Zone: public
  Flow Statistics :
  Flow Input statistics :
    Self packets :                    0
    ICMP packets :                    0
    VPN packets :                     0
    Multicast packets :               0
    Bytes permitted by policy :       0
    Connections established :         0
  Flow Output statistics:
    Multicast packets :               0
```

```
      Bytes permitted by policy :        0
    Flow error statistics (Packets dropped due to):
      Address spoofing:                  0
      Authentication failed:             0
      Incoming NAT errors:               0
      Invalid zone received packet:      0
      Multiple user authentications:     0
      Multiple incoming NAT:             0
      No parent for a gate:              0
      No one interested in self packets: 0
      No minor session:                  0
      No more sessions:                  0
      No NAT gate:                       0
      No route present:                  0
      No SA for incoming SPI:            0
      No tunnel found:                   0
      No session for a gate:             0
      No zone or NULL zone binding       0
      Policy denied:                     0
      Security association not active:   0
      TCP sequence number out of window: 0
      Syn-attack protection:             0
      User authentication errors:        0
    Protocol inet, MTU: 1500, Generation: 150, Route table: 0
      Flags: Sendbcast-pkt-to-re
      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: 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     10.1.65.1/24
ge-0/0/1              up    up
ge-0/0/2              up    up
e-0/0/3               up    up
ge-0/0/4              up    up
```

show interfaces controller (Channelized E1 IQ with Logical E1)

```
user@host> show interfaces controller ce1-1/2/6

Controller                                                      Admin Link
ce1-1/2/6                                                       up    up
e1-1/2/6                                                        up    up
```

show interfaces controller (Channelized E1 IQ with Logical DS0)

```
user@host> show interfaces controller ce1-1/2/3

Controller                                                      Admin Link
ce1-1/2/3                                                       up    up
ds-1/2/3:1                                                      up    up
ds-1/2/3:2                                                      up    up
```

show interfaces descriptions

```
user@host> show interfaces descriptions
 Interface       Admin Link Description
 so-1/0/0        up    up   M20-3#1
 so-2/0/0        up    up   GSR-12#1
 ge-3/0/0        up    up   SMB-OSPF_Area300
 so-3/3/0        up    up   GSR-13#1
 so-3/3/1        up    up   GSR-13#2
 ge-4/0/0        up    up   T320-7#1
 ge-5/0/0        up    up   T320-7#2
 so-7/1/0        up    up   M160-6#1
 ge-8/0/0        up    up   T320-7#3
 ge-9/0/0        up    up   T320-7#4
 so-10/0/0       up    up   M160-6#2
 so-13/0/0       up    up   M20-3#2
 so-14/0/0       up    up   GSR-12#2
 ge-15/0/0       up    up   SMB-OSPF_Area100
 ge-15/0/1       up    up   GSR-13#3
```

show interfaces destination-class all

```
user@host> show interfaces destination-class all
Logical interface so-4/0/0.0
                                          Packets                Bytes
     Destination class           (packet-per-second)    (bits-per-second)
                        gold                        0                    0
                                  (                 0) (                 0)
                      silver                        0                    0
                                  (                 0) (                 0)
  Logical interface so-0/1/3.0
                                          Packets                Bytes
     Destination class           (packet-per-second)    (bits-per-second)
                        gold                        0                    0
                                  (                 0) (                 0)
                      silver                        0                    0
                                  (                 0) (                 0)
```

show interfaces diagnostics optics

```
user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
    Laser bias current                        :  7.408 mA
    Laser output power                        :  0.3500 mW / -4.56 dBm
    Module temperature                        :  23 degrees C / 73 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.0002 mW / -36.99 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
```

```
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  On
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  On
    Laser bias current high alarm threshold   :  17.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6310 mW / -2.00 dBm
    Laser output power low alarm threshold    :  0.0670 mW / -11.74 dBm
    Laser output power high warning threshold :  0.6310 mW / -2.00 dBm
    Laser output power low warning threshold  :  0.0790 mW / -11.02 dBm
    Module temperature high alarm threshold   :  95 degrees C / 203 degrees F
    Module temperature low alarm threshold    :  -25 degrees C / -13 degrees F
    Module temperature high warning threshold :  90 degrees C / 194 degrees F
    Module temperature low warning threshold  :  -20 degrees C / -4 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2590 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7940 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0158 mW / -18.01 dBm
```

show interfaces far-end-interval coc12-5/2/0

```
user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0
```

```
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
  04:00-04:15:
...
```

show interfaces far-end-interval coc1-5/2/1:1

```
user@host> run show interfaces far-end-interval coc1-5/2/1:1
 Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
  05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
  04:00-04:15:
```

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
ge-5/0/0        up    up
ge-5/0/0.0      up    up   any                        f-any
                           inet                       f-inet
                           multiservice
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
```

```
pe-0/3/0        up    up
vt-0/3/0        up    up
at-1/0/0        up    up
at-1/0/0.0      up    up   inet
                           iso
at-1/1/0        up    down
at-1/1/0.0      up    down inet
                           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: 203.0.113.255
```

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
                                          Packets                Bytes
      Source class                (packet-per-second)    (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
             Runt frames                              0
             Jabber frames                            0
             Fragment frames                          0
```

```
              Discarded frames                          0
            Autonegotiation information:
              Negotiation status: Complete
              Link partner:
                  Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link partner
Speed: 100 Mbps
              Local resolution:
                  Flow control: None, Remote fault: Link OK
```

show interfaces transport pm (SRX devices)

```
user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
  14:45-current          Elapse time:900 Seconds
  Near End       Suspect Flag:False             Reason:None
    PM                    COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                    COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  OTU-BBE                  0                800              No              No
  OTU-ES                   0                135              No              No
  OTU-SES                  0                90               No              No
  OTU-UAS                  0                90               No              No
  Near End       Suspect Flag:False             Reason:None
    PM                    COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  427              90               No              No
  Far End        Suspect Flag:True              Reason:Unknown
    PM                    COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  ODU-BBE                  0                800              No              No
  ODU-ES                   0                135              No              No
  ODU-SES                  0                90               No              No
  ODU-UAS                  0                90               No              No
  FEC            Suspect Flag:False             Reason:None
    PM                    COUNT           THRESHOLD       TCA-ENABLED     TCA-RAISED
  FEC-CorrectedErr         2008544300       0                NA              NA
```

```
  FEC-UncorrectedWords    0               0               NA              NA
  BER             Suspect Flag:False              Reason:None
    PM                  MIN       MAX       AVG       THRESHOLD      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 extensive (QFX5130-32CD)

```
user@host> show interfaces et-0/0/29 extensive | no-more
Physical interface: et-0/0/29, Enabled, Physical link is Up
  Interface index: 1086, SNMP ifIndex: 549, Generation: 618475300929
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Media type: Copper
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 0c:59:9c:81:fb:12, Hardware address: 0c:59:9c:81:fb:12
  Last flapped   : 2020-09-14 06:27:45 PDT (1d 01:34 ago)
  Statistics last cleared: Never
  Traffic statistics:
```

```
 Input  bytes  :      130245061850190                   0 bps
 Output bytes  :      132765627331264                   0 bps
 Input  packets:          86830042098                   0 pps
 Output packets:          88510419103                   0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 4, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 5 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets
  0                              0                   0                    0
  3                              0                   0                    0
  4                              0                   0                    0
  7                     88510418098          88510418098                    0
  8                              0                   0                    0
Queue number:         Mapped forwarding classes
  0                   best-effort
  3                   fcoe
  4                   no-loss
  7                   network-control
  8                   mcast
Active alarms  : None
Active defects : None
PCS statistics                      Seconds
  Bit errors                             0
  Errored blocks                         0
Ethernet FEC Mode  :                FEC119
Ethernet FEC statistics             Errors
  FEC Corrected Errors           5796771110
  FEC Uncorrected Errors                 0
  FEC Corrected Errors Rate              0
  FEC Uncorrected Errors Rate            0
MAC statistics:                     Receive         Transmit
  Total octets               130245061850190  132765627331264
  Total packets                  86830042098      88510419103
  Unicast packets                86830041265      88510418256
  Broadcast packets                      179              190
  Multicast packets                      654              657
  CRC/Align errors                         0                0
  FIFO errors                              0                0
```

```
    MAC control frames                      0               0
    MAC pause frames                        0               0
    Oversized frames                        0
    Jabber frames                           0
    Fragment frames                         0
    VLAN tagged frames                      0
    Code violations                         0
    Total errors                            0               0
  MAC Priority Flow Control Statistics:
    Priority :  0                           0               0
    Priority :  1                           0               0
    Priority :  2                           0               0
    Priority :  3                           0               0
    Priority :  4                           0               0
    Priority :  5                           0               0
    Priority :  6                           0               0
    Priority :  7                           0               0
  Filter statistics:
    Input packet count                      0
    Input packet rejects                    0
    Input DA rejects                        0
    Input SA rejects                        0
    Output packet count                                     0
    Output packet pad count                                 0
    Output packet error count                               0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot:  ( )
  CoS information:
    Direction :
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                 :  Disable
```

show interfaces (PTX10001-36MR) (400G ZR and 400G ZR-M optics)

```
user@host> show interfaces et-0/0/10
Physical interface: et-0/0/10, Enabled, Physical link is Up
  Interface index: 1016, SNMP ifIndex: 519
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps,
  BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
```

```
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  Optic-loopback : Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 40:de:ad:8c:ca:22, Hardware address: 40:de:ad:8c:ca:22
  Last flapped   : 2022-02-28 07:32:26 UTC (03:58:05 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC Mode  :                FEC119
  Ethernet FEC statistics             Errors
    FEC Corrected Errors              179957
    FEC Uncorrected Errors                 0
    FEC Corrected Errors Rate             18
    FEC Uncorrected Errors Rate            0
  Optic FEC Mode  :                     OFEC
  Optic FEC statistics:
    Corrected Errors                          2695089428427
    Uncorrected Words                                     0
    Corrected Error rate                          194566696
    Uncorrected Error rate                                0
    Corrected Error Ratio (      14407 seconds average)  4.08e-04
  PRBS Mode : Disabled
  Interface transmit statistics: Disabled
  Link Degrade :
    Link Monitoring                   :  Disable
```

show interfaces (PTX10001-36MR) (Junos OS Evolved Release)

```
user@host> show interfaces et-0/1/0 media
Physical interface: et-0/1/0, Enabled, Physical link is Down
  Interface index: 1017, SNMP ifIndex: 554
```

```
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Media type: Copper
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 88:d9:8f:08:46:57, Hardware address: 88:d9:8f:08:46:57
  Last flapped   : 2021-08-30 02:45:11 PDT (00:00:34 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK, REMOTE-FAULT
  Ethernet FEC Mode  :                FEC119
  Ethernet FEC statistics             Errors
    FEC Corrected Errors            545479262
    FEC Uncorrected Errors                 24
    FEC Corrected Errors Rate            7942
    FEC Uncorrected Errors Rate             0
  MAC statistics:
    Input bytes: 0, Input packets: 0, Output bytes: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  PRBS Mode : Enabled
  PRBS Pattern : 7
  PRBS Statistics
    Lane 0 : Number of iterations in error :   4538027 Total iterations :     4538027 Monitored
Seconds: 2
    Lane 1 : Number of iterations in error :   2269036 Total iterations :     2269036 Monitored
Seconds: 1
    Lane 2 : Number of iterations in error :   2269247 Total iterations :     2269247 Monitored
Seconds: 1
    Lane 3 : Number of iterations in error :   4538551 Total iterations :     4538551 Monitored
Seconds: 2
    Lane 4 : Number of iterations in error :   4537975 Total iterations :     4537975 Monitored
Seconds: 2
    Lane 5 : Number of iterations in error :   4538053 Total iterations :     4538053 Monitored
Seconds: 2
    Lane 6 : Number of iterations in error :   4538555 Total iterations :     4538555 Monitored
Seconds: 2
    Lane 7 : Number of iterations in error :   4538433 Total iterations :     4538433 Monitored
Seconds: 2
  Interface transmit statistics: Disabled
  Link Degrade :
```

```
    Link Monitoring                   :  Disable

  Logical interface et-0/1/0.16386 (Index 7003) (SNMP ifIndex 560)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

show interfaces extensive (ACX5448, ACX5448-D, ACX5448-M and ACX710)

```
user@host> show interfaces xe-0/0/18 extensive
Physical interface: xe-0/0/18, Enabled, Physical link is Up
Interface index: 154, SNMP ifIndex: 525, Generation: 185
Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 1000mbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
Auto-negotiation: Enabled, Remote fault: Offline,
Media type: Fiber
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Schedulers : 0
Hold-times : Up 0 ms, Down 0 ms
Damping : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state: unsuppressed
Current address: 98:a4:04:79:50:14, Hardware address: 98:a4:04:79:50:14
Last flapped : 2021-06-29 09:33:22 PDT (13:43:07 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 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
Ethernet FEC statistics Errors
FEC Corrected Errors 0
FEC Uncorrected Errors 0
FEC Corrected Errors Rate 0
FEC Uncorrected Errors Rate 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
```

```
PRBS Mode : Disabled
Filter statistics:
Input packet count 0
Input packet rejects 0
Input DA rejects 0
Input SA rejects 0
Output packet count 0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK, Link partner
Speed: 1000 Mbps
Local resolution:
Flow control: Symmetric, Remote fault: Link OK, Local link Speed: 1000 Mbps, Link mode: Full-
duplex
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 950000000 95 0 low none
3 network-control 5 50000000 5 0 low none
Interface transmit statistics: Disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

Command modified in Junos OS Release 9.5 for SRX Series devices.

Command modified in Junos OS Release 19.3R1 for MX Series Routers.

Starting in Junos OS Release 19.3R1, Output fields `Ifindex` and `speed` is modified in the `show interfaces` *`interface name`* `extensive` command, on all MX Series routers.

- The default behavior of WAN-PHY interface remains the same.The new `precise-bandwidth` option reflects the new speed (9.294-Gbps) configured on the supported line cards.

- The WAN-PHY framing mode is supported only on MPC5E and MPC6E line cards.

Starting in Junos OS Release 19.3R1, class of service (CoS) features can be configured on the physical interface with speed rates of 1-Gbps, 10-Gbps, 40-Gbps, and 100-Gbps to provide better bandwidth for processing traffic during congestion using variant speeds.

Release History Table

Release	Description
20.4R1	Starting in Junos OS Release 20.4R1, we support G.fast and Annex J specification with SFP xDSL for ADSL2/ADSL2+ and all VDSL2 profiles on SRX380, SRX300, SRX320, SRX340, and SRX345 devices.
19.2R3	In Junos OS Releases 19.2R3, 19.3R3, 19.4R3, 20.1R2, and 20.2R1, on QFX5120-48Y switch, the `show interfaces` *`interface-name`*`<media><extensive>` command displays the autonegotiation status only for the interface that supports autonegotiation.
18.4R1	Starting in Junos OS Release 18.4R1, Output fields `Next-hop` and `vpls-status` is displayed in the `show interfaces` *`interface name`* `detail` command, only for Layer 2 protocols on MX480 routers.

RELATED DOCUMENTATION

Understanding Layer 2 Interfaces on Security Devices

Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration

Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers

dsl-sfp-options

show interfaces diagnostics optics

IN THIS SECTION

- Syntax | **1024**
- Description | **1024**
- Options | **1024**
- Required Privilege Level | **1025**

Output Fields | 1025
Sample Output | 1029
Release Information | 1030

Syntax

```
show interfaces diagnostics optics interface-name
```

Description

Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP) installed in SRX Series Services Gateways. 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.

NOTE: In a chassis cluster, the `show interfaces diagnostics optics` command works only on the node that is primary in redundancy group 0 (RG0).

Options

interface-name Name of the interface associated with the port in which the transceiver is installed: `ge-`*`fpc/pic/port`*.

Required Privilege Level

view

Output Fields

Table 65 on page 1025 lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the general order in which they appear.

Table 65: show interfaces diagnostics optics Output Fields

Field Name	Field Description
`Physical interface`	Displays the name of the physical interface.
`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`	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 `On` or `Off`.
`Laser bias current low alarm`	Displays whether the laser bias power setting low alarm is `On` or `Off`.
`Laser bias current high warning`	Displays whether the laser bias power setting high warning is `On` or `Off`.

Table 65: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser bias current low warning`	Displays whether the laser bias power setting low warning is `On` or `Off`.
`Laser output power high alarm`	Displays whether the laser output power high alarm is `On` or `Off`.
`Laser output power low alarm`	Displays whether the laser output power low alarm is `On` or `Off`.
`Laser output power high warning`	Displays whether the laser output power high warning is `On` or `Off`.
`Laser output power low warning`	Displays whether the laser output power low warning is `On` or `Off`.
`Module temperature high alarm`	Displays whether the module temperature high alarm is `On` or `Off`.
`Module temperature low alarm`	Displays whether the module temperature low alarm is `On` or `Off`.
`Module temperature high warning`	Displays whether the module temperature high warning is `On` or `Off`.
`Module temperature low warning`	Displays whether the module temperature low warning is `On` or `Off`.
`Module voltage high alarm`	Displays whether the module voltage high alarm is `On` or `Off`.
`Module voltage low alarm`	Displays whether the module voltage low alarm is `On` or `Off`.
`Module voltage high warning`	Displays whether the module voltage high warning is `On` or `Off`.
`Module voltage low warning`	Displays whether the module voltage low warning is `On` or `Off`.
`Laser rx power high alarm`	Displays whether the receive laser power high alarm is `On` or `Off`.
`Laser rx power low alarm`	Displays whether the receive laser power low alarm is `On` or `Off`.

Table 65: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
`Laser rx power high warning`	Displays whether the receive laser power high warning is `On` or `Off`.
`Laser rx power low warning`	Displays whether the receive laser power low warning is `On` or `Off`.
`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.
`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.

Table 65: show interfaces diagnostics optics Output Fields *(Continued)*

Field Name	Field Description
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.
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 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
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces | 2

show interfaces flow-statistics

IN THIS SECTION

- Syntax | 1031
- Description | 1031
- Options | 1031
- Required Privilege Level | 1032
- Output Fields | 1032
- Sample Output | 1035
- Release Information | 1037

Syntax

```
show interfaces flow-statistics                    <interface-name>
```

Description

Display interfaces flow statistics.

Options

Interface-name —(Optional) Display flow statistics about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and *port* with the port number. For a complete list, see the "Interface Naming Conventions" on page 9.

- at-*pim*/0/*port*—ATM-over-ADSL or ATM-over-SHDSL interface.
- br-*pim*/0/*port*—Basic Rate Interface for establishing ISDN connections.
- ce1-*pim*/0/*port*—Channelized E1 interface.
- 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).

Required Privilege Level

view

Output Fields

Table 66 on page 1032 lists the output fields for the `show interfaces flow-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 66: show interfaces flow-statistics Output Fields

Field Name	Field Description
`Traffic statistics`	Number of packets and bytes transmitted and received on the physical interface.
`Local statistics`	Number of packets and bytes transmitted and received on the physical interface.
`Transit statistics`	Number of packets and bytes transiting the physical interface.
`Flow input statistics`	Statistics on packets received by flow module.
`Flow output statistics`	Statistics on packets sent by flow module.
`Flow error statistics`	Packet drop statistics for the flow module. For further details, see Table 67 on page 1032.

Table 67: Flow Error Statistics (Packet Drop Statistics for the Flow Module)

Error	Error Description
Screen:	
Address spoofing	The packet was dropped when the screen module detected address spoofing.

Syn-attack protection	The packet was dropped because of SYN attack protection or SYN cookie protection.
VPN:	
Authentication failed	The packet was dropped because the IPsec Encapsulating Security Payload (ESP) or Authentication Header (AH) authentication failed.
No SA for incoming SPI	The packet was dropped because the incoming IPsec packet's security parameter index (SPI) does not match any known SPI.
Security association not active	The packet was dropped because an IPsec packet was received for an inactive SA.
NAT:	
Incoming NAT errors	The source NAT rule search failed, an invalid source NAT binding was found, or the NAT allocation failed.
Multiple incoming NAT	Sometimes packets are looped through the system more than once; if source NAT is specified more than once, the packet will be dropped.
Auth:	
Multiple user authentications	Sometimes packets are looped through the system more than once. Each time a packet passes through the system, that packet must be permitted by a policy. If the packet matches more than one policy that specifies user authentication, then it will be dropped.
User authentication errors	Packet was dropped because policy requires authentication; however: • Only Telnet, FTP, and HTTP traffic can be authenticated. • The corresponding authentication entry could not be found, if web-auth is specified. • The maximum number of authenticated sessions per user was exceeded.
Flow:	

No one interested in self packets	This counter is incremented for one of the following reasons: • The outbound interface is a self interface, but the packet is not marked as a to-self packet and the destination address is in a source NAT pool. • No service is interested in the to-self packet • When a zone has ident-reset service enabled, the TCP RST to IDENT request for port 113 is sent back and this counter is incremented.
No minor session	The packet was dropped because no minor sessions are available and a minor session was requested. Minor sessions are allocated for storing additional TCP state information.
No more sessions	The packet was dropped because there were no more free sessions available.
No route present	The packet was dropped because a valid route was not available to forward the packet. For new sessions, the counter is incremented for one of the following reasons: • No valid route was found to forward the packet. • A discard or reject route was found. • The route could not be added due to lack of memory. • The reverse path forwarding check failed for an incoming multicast packet. For existing sessions, the prior route was changed or deleted, or a more specific route was added. The session is rerouted, and this reroute could fail because: • A new route could not be found; either the previous route was removed, or the route was changed to discard or reject. • Multiple packets may concurrently force rerouting to occur, and only one packet can successfully complete the rerouting process. Other packets will be dropped. • The route table was locked for updates by the Routing Engine. Packets that match a new session are retried, whereas packets that match an existing session are not.
No tunnel found	The packet was dropped because a valid tunnel could not be found

No session for a gate	This counter is incremented when a packet is destined for an ALG, and the ALG decides to drop this packet.
No zone or NULL zone binding	The packet was dropped because its incoming interface was not bound to any zone.
Policy denied	The error counter is incremented for one of the following reasons: • Source and/or destination NAT has occurred and policy says to drop the packet. • Policy specifies user authentication, which failed. • Policy was configured to deny this packet.
TCP sequence number out of window	A TCP packet with a sequence number failed the TCP sequence number check that was received.
Counters Not Currently in Use	
No parent for a gate	-
Invalid zone received packet	-
No NAT gate	-

Sample Output

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 igmp 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
```

Release Information

Command introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

Understanding Traffic Processing on Security Devices

Understanding Interfaces | 2

show interfaces queue

IN THIS SECTION

- Syntax | 1037
- Description | 1038
- Options | 1038
- Required Privilege Level | 1038
- Output Fields | 1039
- Sample Output | 1041
- Sample Output | 1043
- Release Information | 1044

Syntax

```
show interfaces queue
<both-ingress-egress>
<egress>
<forwarding-class forwarding-class>
<ingress>
```

```
<interface-name interface-name>
<l2-statistics>
```

Description

Display class-of-service (CoS) queue information for physical interfaces.

NOTE: The queue depth information is only available on vSRX and SRX1500, SRX4100, SRX4200 and SRX4600 platforms.

Options

none	Show detailed CoS queue statistics for all physical interfaces.
both-ingress-egress	Display both ingress and egress queue statistics.
egress	Display egress queue statistics.
forwarding-class *forwarding-class*	(Optional) Forwarding class name for this queue. Show detailed CoS statistics for the queue that is associated with the specified forwarding class.
ingress	Display ingress queue statistics.
interface-name *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.

Required Privilege Level

view

Output Fields

Table 68 on page 1039 lists the output fields for the `show interfaces queue` command. Output fields are listed in the approximate order in which they appear.

Table 68: 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*.
`Interface index`	Index number of the physical interface. The number reflects the interface’s initialization sequence.
`SNMP ifIndex`	SNMP index number for the interface.
`Forwarding classes supported`	Total number of forwarding classes supported on the specified interface.
`Forwarding classes in use`	Total number of forwarding classes 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.
The following output fields are applicable to both the interface component and Packet Forwarding Engine component in the `show interfaces queue` command:	
`Queue`	Queue number.
`Forwarding classes`	Forwarding class name.

Table 68: show interfaces queue Output Fields *(Continued)*

Field Name	Field Description
Queued Packets	Number of packets in this queue.
Queued Bytes	Number of bytes in this queue.
Transmitted Packets	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.
Transmitted Bytes	Number of bytes transmitted by this queue.
Tail-dropped packets	Number of packets dropped because of tail drop.
RL-dropped bytes	Number of bytes dropped because of rate limiting.
RED-dropped packets	Number of packets dropped because of random early detection (RED).
RED-dropped bytes	Number of bytes dropped because of RED. • 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.
Queue Buffer Usage:	• Reserved buffer—The size of the memory buffer that is allocated for storing packets • Current—The amount of buffer memory that is currently in use on this queue.
Queue-Depth	Current—The maximum number of bytes in this queue, that is currently in use on this queue.

Sample Output

show interfaces queue (vSRX)

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: 135, SNMP ifIndex: 510
Forwarding classes: 8 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :                14686                    0 pps
    Bytes                :               616812                    0 bps
  Transmitted:
    Packets              :                14686                    0 pps
    Bytes                :               616812                    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: 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 Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
    Current              :                    0
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 Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
```

```
    Current             :                   0
...
```

Sample Output

show interfaces queue (vSRX)

```
user@host> show interfaces queue ge-0/0/3 forwarding-class ef
Physical interface: ge-0/0/3, Enabled, Physical link is Up
  Interface index: 143, SNMP ifIndex: 510
Forwarding classes: 8 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 1, Forwarding classes: ef
  Queued:
    Packets              :             55034875               885424 pps
    Bytes                :         1526912538034            835840256 bps
  Transmitted:
    Packets              :            2772633112                    0 pps
    Bytes                :         1512013543328                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RL-dropped packets   :                    0                    0 pps
    RL-dropped bytes     :                    0                    0 bps
    RED-dropped packets  :            126262505               885424 pps
     Low                 :                    0                    0 pps
     Medium-low          :                    0                    0 pps
     Medium-high         :                    0                    0 pps
     High                :            126262505               885424 pps
    RED-dropped bytes    :          14898975590            835840728 bps
     Low                 :                    0                    0 bps
     Medium-low          :                    0                    0 bps
     Medium-high         :                    0                    0 bps
     High                :          14898975590            835840728 bps
  Queue Buffer Usage:
    Reserved buffer      :                 9192 bytes
  Queue-depth bytes      :
    Current              :                18998
```

Release Information

Command introduced in Junos OS Release 15.1X49-D30.

RELATED DOCUMENTATION

Understanding Class of Service

show interfaces statistics (View)

IN THIS SECTION

- Syntax | 1044
- Description | 1044
- Required Privilege Level | 1045
- Sample Output | 1045
- Release Information | 1045

Syntax

```
show interfaces statistics interface-name
```

Description

Displays the interface input and output statistics for physical and logical interface.

Required Privilege Level

view

Sample Output

show interfaces statistics

```
user@host> show interfaces statistics st0.1
Logical interface st0.1 (Index 91) (SNMP ifIndex 268)
     Flags: Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
     Input packets : 2743333
     Output packets: 6790470992
     Security: Zone: untrust
     Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp nhrp ospf pgm pim rip
router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http https ike netconf ping
reverse-telnet
     reverse-ssh rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text xnm-ssl
lsping ntp sip
     Protocol inet, MTU: 9192
       Addresses, Flags: Is-Preferred Is-Primary
         Destination: 192.167.1.0/30, Local: 192.167.1.1
```

Release Information

Command introduced in Junos OS Release 10.1.

RELATED DOCUMENTATION

Understanding Interfaces | 2

show interfaces terse zone

IN THIS SECTION

- Syntax | 1046
- Description | 1046
- Options | 1046
- Required Privilege Level | 1046
- Sample Output | 1047
- Release Information | 1047

Syntax

```
show interfaces terse zone
```

Description

Display summary information about zone interfaces.

Options

This command has no options.

Required Privilege Level

view

Sample Output

show interface terse zone

```
user@host> show interface terse zone
Interface     Admin    Link   Proto     Local           Remote     Zone
ge-0/0/0.0      up      up     inet     1.4.253.251/16             trust
```

Release Information

Command introduced in Junos OS Release 12.3X48-D20.

show ipv6 neighbors

IN THIS SECTION

- Syntax | 1047
- Description | 1048
- Options | 1048
- Required Privilege Level | 1048
- Output Fields | 1048
- Sample Output | 1049
- Release Information | 1049

Syntax

```
show ipv6 neighbors
```

Description

Display information about the IPv6 neighbor cache.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 69 on page 1048 lists the output fields for the `show ipv6 neighbors` command. Output fields are listed in the approximate order in which they appear.

Table 69: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up, down, incomplete, reachable, stale, or unreachable.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no.

Table 69: show ipv6 neighbors Output Fields *(Continued)*

Field Name	Field Description
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no.
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address     Linklayer Address      State       Exp Rtr Secure     Interface
            10:1::2          00:00:0a:00:00:00      reachable   17   yes     no     reth0.0
            11:11::2         00:19:e2:4b:61:83      stale       1197 yes     no     at-1/0/0.0
            12:12::2         00:19:e2:4b:61:83      stale       1188 yes     no     at-3/0/0.0
```

Release Information

Command introduced in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

clear ipv6 neighbors | 781

show lacp interfaces (View)

IN THIS SECTION

- Syntax | 1050
- Description | 1050
- Options | 1050
- Required Privilege Level | 1051
- Output Fields | 1051
- Sample Output | 1054
- Release Information | 1056

Syntax

```
show lacp interfaces interface-name
```

Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet interface, redundant Ethernet interface, Gigabit Ethernet interface, or 10-Gigabit Ethernet interface. If you do not specify an interface name, LACP information for all interfaces is displayed.

Options

none	Display LACP information for all interfaces.
interface-name	(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—ae*number*
- Redundant Ethernet—reth*number*
- Gigabit Ethernet—ge-*fpc*/*pic*/*port*
- 10-Gigabit Ethernet—xe-*fpc*/*pic*/*port*

NOTE: The `show lacp interfaces` command returns the following error message if your system is not configured in either active or passive LACP mode:
"Warning: lacp subsystem not running – not needed by configuration"

Required Privilege Level

view

Output Fields

Table 70 on page 1051 lists the output fields for the `show lacp interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 70: show lacp interfaces Output Fields

Field Name	Field Description
`Aggregated interface`	Aggregated interface value.

Table 70: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description

LACP State	LACP state information for each aggregated interface:
	• Role—Role played by the interface. It can be one of the following:
	  • 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 a link aggregation control protocol data unit (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 link aggregation control PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (Long Timeout or Short Timeout).
	• Activity—Actor or partner's port activity. Passive indicates the port's preference for not transmitting link aggregation control 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.

Table 70: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description

`LACP Protocol`	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• `Receive State`—One of the following values:
	  • `Current`—The state machine receives a link aggregation control PDU and enters the `Current` state.
	  • `Defaulted`—If no link aggregation control PDU is received before the timer for the `Current` state expires a second time, the state machine enters the `Defaulted` state.
	  • `Expired`—If no link aggregation control 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:
	  • `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:
	  • `Attached`—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.

Table 70: show lacp interfaces Output Fields *(Continued)*

Field Name	Field Description
	• 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. • Waiting—Multiplexer state machine is in a holding process, awaiting an outcome.

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-2/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-2/2/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-2/0/0                  Current   Fast periodic Collecting distributing
      ge-2/0/1                  Current   Fast periodic Collecting distributing
      ge-2/2/0                  Current   Fast periodic Collecting distributing
      ge-2/2/1                  Current   Fast periodic Collecting distributing
```

show lacp interfaces (Redundant Ethernet)

```
user@host> show lacp interfaces reth0
Aggregated interface: reth0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-11/0/0      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/0    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/1    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/2    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3      Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-11/0/3    Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/0     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-3/0/3     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-11/0/0                Current   Fast periodic Collecting distributing
      ge-11/0/1                Current   Fast periodic Collecting distributing
      ge-11/0/2                Current   Fast periodic Collecting distributing
      ge-11/0/3                Current   Fast periodic Collecting distributing
      ge-3/0/0                 Current   Fast periodic Collecting distributing
      ge-3/0/1                 Current   Fast periodic Collecting distributing
      ge-3/0/2                 Current   Fast periodic Collecting distributing
      ge-3/0/3                 Current   Fast periodic Collecting distributing
{primary:node1}
```

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
```

Release Information

Command modified in Junos OS Release 10.2.

RELATED DOCUMENTATION

- Verifying LACP on Redundant Ethernet Interfaces

show lacp statistics interfaces (View)

IN THIS SECTION

- Syntax | 1056
- Description | 1057
- Options | 1057
- Required Privilege Level | 1057
- Output Fields | 1057
- Sample Output | 1058
- Release Information | 1058

Syntax

```
show lacp statistics interfaces interface-name
```

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

Output Fields

Table 71 on page 1057 lists the output fields for the `show lacp statistics interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 71: show lacp statistics interfaces Output Fields

Field Name	Field Description
`Aggregated interface`	Aggregated interface value.

Table 71: show lacp statistics interfaces Output Fields *(Continued)*

Field Name	Field Description
LACP Statistics	LACP statistics provide the following information: • 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. **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.

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
```

Release Information

Command modified in Release 10.2 of Junos OS.

Command introduced in Release 11.1 of Junos OS.

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*
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*

show modem wireless firmware

IN THIS SECTION

- Syntax | **1059**
- Description | **1059**
- Options | **1060**
- Required Privilege Level | **1060**
- Output Fields | **1060**
- Sample Output | **1062**
- Release Information | **1063**

Syntax

```
show modem wireless firmware interface-name
```

Description

Display modem firmware details for the LTE Mini-PIM.

Options

- *interface-name*—The LTE interface is `cl-x/0/0`, where *x* is the slot number in which the LTE Mini-PIM is installed.

Required Privilege Level

view

Output Fields

Table 72 on page 1060 lists some of the output fields for the `show modem wireless firmware` command. Output fields are listed in the approximate order in which they appear.

Table 72: show modem wireless firmware Output Fields

Field Name	Description
`LTE mPIM firmware details`	Displays the details of the firmware installed on the LTE Mini-PIM.
`Wireless modem firmware details`	Displays the details of the modem firmware.
`OTA status`	Displays the status of over-the-air (OTA) upgrade. The OTA upgrade can be enabled or disabled on the LTE Mini-PIM. OTA upgrade is disabled by default.

Table 72: show modem wireless firmware Output Fields *(Continued)*

Field Name	Description

Status of SIM	• Number of SIM—Number of SIM cards installed.
	• Slot of active—The slot in which the active SIM card is installed.
	• SIM state—Indicates whether the SIM card is present in the slot.
	• Modem PIN security status—Indicates the security status of the SIM. If the SIM is locked by using the request modem wireless sim-lock enable command, then the security status is displayed as enabled.
	• SIM status—Status of the Subscriber Identity Module (SIM) in the LTE Mini-PIM. The status can be one of the following:
	  • SIM Okay
	  • No status—The device is being powered on or powered off, or the SIM card has been removed from the slot.
	  • SIM init failure—There is a problem with the SIM; the SIM might need to be replaced.
	  • SIM locked
	  • PIN1 blocked—Obtain a PIN unblocking key (PUK) to unblock the SIM.
	  • PIN1 rejected—The wrong PIN was entered.
	  • PIN2 rejected—The wrong PIN was entered.
	  • Network rejected
	• SIM user operation needed—Action required by the user. This can be one of the following:
	  • No op—No user operation required.
	  • Enter PIN—Enter the personal identification number (PIN) to unlock the SIM card.
	  • Enter PUK—Enter the PUK to unblock the SIM card.
	• Retries remaining—If the value of SIM user operation needed is Enter PIN, this is the number of PIN unlock attempts remaining before the modem is blocked. If the PIN is entered incorrectly three consecutive times, the SIM card is blocked.

Table 72: show modem wireless firmware Output Fields *(Continued)*

Field Name	Description
	If the value of `SIM user operation needed` is `Enter PUK`, this is the number of unblock attempts remaining before the modem is unusable. If the PUK is entered incorrectly ten times, the SIM card must be returned to the service provider for reactivation.

Sample Output

show modem wireless firmware

```
user@host> show modem wireless firmware cl-1/0/0
 LTE mPIM firmware details
  Product name: Junos LTE mPIM
  Serial number: AG50071852
  Hardware version: AcceleratedConcepts/sprite
  Firmware version: 17.4.3
  MAC: 00:00:5e:00:a0:61
  System uptime: 3430 seconds
Wireless modem firmware details
  Modem firmware version: 9999999_9904609_SWI9X30C_02.23.00.00_00_GENERIC_002.018_000
  Modem Firmware build date: 22/10/2016
  Card type: MC7430
  Modem manufacturer: Sierra Wireless, Inc
  Hardware version: 1.0
  Power & Temperature: Normal 3343 mV, Normal 30.00 C
OTA status
  State: Enabled
  New firmware available: No
Number of SIM: 2
Slot of active: 2
Status of SIM 1
  SIM state: SIM present
  Modem PIN security status: Disabled
  SIM status: SIM Okay
  SIM user operation needed: No Op
  Retries remaining: 3
```

```
Status of SIM 2
  SIM state: SIM present
  Modem PIN security status: Disabled
  SIM status: SIM Okay
  SIM user operation needed: No Op
  Retries remaining: 3
```

Release Information

Command introduced in Junos OS 15.1X49-D100

RELATED DOCUMENTATION

show modem wireless network | 1063

show modem wireless network

IN THIS SECTION

- Syntax | 1064
- Description | 1064
- Options | 1064
- Required Privilege Level | 1064
- Output Fields | 1064
- Sample Output | 1067
- Release Information | 1068

Syntax

```
show modem wireless network interface-name
```

Description

Display the status of the modem and the status of the network connection for the LTE Mini-PIM.

Options

- *interface-name*—The LTE interface is `cl-x/0/0`, where *x* is the slot number in which the LTE Mini-PIM is installed.

Required Privilege Level

view

Output Fields

Table 73 on page 1065 lists some of the output fields for the `show modem wireless network` command. Output fields are listed in the approximate order in which they appear.

Table 73: show modem wireless network Output Fields

Field Name	Field Description
`Current Modem Status`	Status of the modem on the Mini-PIM. The status can be one of the following states: • Disconnected • Calling • Connected
`Current Service Status`	Status of the network connection. The status can be one of the following states: • Normal • Emergency Call Only • No Service Available • Unable To Register • Forbidden PLMN • Forbidden Area • Roaming Not Permitted • Account Not Permitted • Modem Not Permitted • Unknown IMSI • Authentication Failure

Table 73: show modem wireless network Output Fields *(Continued)*

Field Name	Field Description
`Current Service Type`	One of the following: • Circuit switched (CS) • Packet switched (PS) • Combo (CS, PS) • Invalid
`Current Service Mode`	One of the following: • Unknown • LTE • DC-HSPA+ • HSPA+ • HSPA • UMTS
`Current Band`	Current radio band in use.
`Mobile Country Code (MCC)`	Number that uniquely identifies the country.
`Mobile Network Code`	Number that uniquely identifies a network within a country.

Sample Output

show modem wireless network

```
user@host> show modem wireless network cl-1/0/0
LTE Connection details
Connected time: 147
IP: 172.16.52.4
Gateway: 172.16.52.5
DNS: 123.123.123.123
Input bps: 0
Output bps: 0
Bytes Received: 1308
Bytes Transferred: 1164
Packets Received: 10
Packets Transferred: 10
Wireless Modem Network Info
Current Modem Status: Connected
Current Service Status: Normal
Current Service Type: PS
Current Service Mode: LTE
Current Band: B3
Network: UNICOM
Mobile Country Code (MCC): 460
Mobile Network Code (MNC): 1
Location Area Code (LAC): 65534
Routing Area Code (RAC): 0
Cell Identification: 4865903
Access Point Name (APN): abcde
Public Land Mobile Network (PLMN): CHN-UNICOM
Physical Cell ID (PCI): 333
International Mobile Subscriber Identification (IMSI): ***************
International Mobile Equipment Identification (IMEI/MEID): ***************
Integrate Circuit Card Identity (ICCID): 89860114721100697502
Reference Signal Receiving Power (RSRP): -97
Reference Signal Receiving Quality (RSRQ): -16
Signal to Interference-plus-Noise Ratio (SiNR): 0
Signal Noise Ratio (SNR): 0
Energy per Chip to Interference (ECIO): 0
```

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless profiles | 1068

show modem wireless firmware | 1059

show modem wireless profiles

IN THIS SECTION

- Syntax | 1068
- Description | 1068
- Options | 1069
- Required Privilege Level | 1069
- Output Fields | 1069
- Sample Output | 1070
- Release Information | 1070

Syntax

```
show modem wireless profiles interface-name slot slot-number
```

Description

Display the profiles configured on the LTE Mini-PIM.

Options

- *interface-name*—The LTE interface is cl-*x*/0/0, where *x* is the slot number in which the LTE Mini-PIM is installed.
- *slot-number*—The slot in which the SIM card is inserted. The value can be either 1 or 2.

Required Privilege Level

view

Output Fields

Table 74 on page 1069 lists some of the output fields for the show modem wireless profiles command. Output fields are listed in the approximate order in which they appear.

Table 74: show modem wireless profiles Output Fields

Field Name	Field Description
Max profiles	The maximum number of profiles available for each SIM card. This value is always 16. The LTE Mini-PIM supports two SIM cards and so you can configure a total of 32 profiles, although only one profile can be active at a time.
Default profile Id	The profile used to connect to the network when there is no profile selected. The default profile ID is always 1.
Profile details	• Username—The username provided by the service provider. • Password—The password provided by the service provider. • Access point name (APN)—The APN provided by the service provider. • Authentication—The protocol used for authentication.

Sample Output

show modem wireless profiles

```
user@host> show modem wireless profiles cl-1/0/0 slot 1
Profile details
  Max profiles: 16
  Default profile Id: 1

Profile 1: ACTIVE
  Valid: TRUE
  Access point name (APN): ctnet
  Authentication: None
Profile 2: Inactive
  Valid: TRUE
  Username: myuser
  Password: 123456
  Access point name (APN): testapn
  Authentication: PAP
Profile 3: Invalid
Profile 4: Invalid
Profile 5: Invalid
Profile 6: Invalid
Profile 7: Invalid
Profile 8: Invalid
Profile 9: Invalid
Profile 10: Invalid
Profile 11: Invalid
Profile 12: Invalid
Profile 13: Invalid
Profile 14: Invalid
Profile 15: Invalid
Profile 16: Invalid
```

Release Information

Command introduced in Junos OS Release 15.1X49-D100.

RELATED DOCUMENTATION

show modem wireless firmware | 1059

show modem wireless network | 1063

show oam ethernet link-fault-management

IN THIS SECTION

- Syntax | 1071
- Description | 1071
- Options | 1072
- Required Privilege Level | 1072
- Output Fields | 1072
- Sample Output | 1078
- Release Information | 1079

Syntax

```
show oam ethernet link-fault-management
<brief | detail>
<interface-name>
```

Description

Display Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.

Options

brief | detail (Optional) Display the specified level of output.

interface-name (Optional) Display link fault management information for the specified Ethernet interface only.

Required Privilege Level

view

Output Fields

Table 75 on page 1072 lists the output fields for the `show oam ethernet link-fault-management` command. Output fields are listed in the approximate order in which they appear.

Table 75: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
`Status`	Status of the established link. • `Fail`—A link fault condition exists. • `Running`—A link fault condition does not exist.	All levels
`Discovery state`	State of the discovery mechanism: • `Passive Wait` • `Send Any` • `Send Local Remote` • `Send Local Remote Ok`	All levels
`Peer address`	Address of the OAM peer.	All levels

Table 75: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Flags`	Information about the interface. • `Remote-Stable`—Indicates remote OAM client acknowledgment of, and satisfaction with, local OAM state information. `False` indicates that remote DTE has either not seen or is unsatisfied with local state information. `True` indicates that remote DTE has seen and is satisfied with local state information. • `Local-Stable`—Indicates local OAM client acknowledgment of, and satisfaction with, remote OAM state information. `False` indicates that local DTE either has not seen or is unsatisfied with remote state information. `True` indicates that local DTE has seen and is satisfied with remote state information. • `Remote-State-Valid`—Indicates the OAM client has received remote state information found within local information TLVs (type, length, values) of received Information OAM PDUs. `False` indicates that the OAM client has not seen remote state information. `True` indicates that the OAM client has seen remote state information.	All levels
`Remote loopback status`	An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels

Table 75: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Remote entity information`	Remote entity information. • `Remote MUX action`—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • `Remote parser action`—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to the higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • `Discovery mode`—Indicates whether discovery mode is active or inactive. • `Unidirectional mode`—Indicates the ability to operate a link in unidirectional mode for diagnostic purposes. • `Remote loopback mode`—Indicates whether remote loopback is supported or not supported. • `Link events`—Indicates whether interpreting link events is supported or not supported on the remote peer. • `Variable requests`—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.	All levels
OAM Receive Statistics		
`Information`	Number of information PDUs received.	`detail`
`Event`	Number of loopback control PDUs received.	`detail`
`Variable request`	Number of variable request PDUs received.	`detail`
`Variable response`	Number of variable response PDUs received.	`detail`
`Loopback control`	Number of loopback control PDUs received.	`detail`

Table 75: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
`Organization specific`	Number of vendor organization specific PDUs received.	`detail`
OAM Transmit Statistics		
`Information`	Number of information PDUs transmitted.	`detail`
`Event`	Number of event notification PDUs transmitted.	`detail`
`Variable request`	Number of variable request PDUs transmitted.	`detail`
`Variable response`	Number of variable response PDUs transmitted.	`detail`
`Loopback control`	Number of loopback control PDUs transmitted.	`detail`
`Organization specific`	Number of vendor organization specific PDUs transmitted.	`detail`
OAM Received Symbol Error Event information		
`Events`	Number of symbol error event TLVs that have been received after the OAM sublayer was reset.	`detail`
`Window`	Symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	`detail`
`Threshold`	Number of errored symbols in the period required for the event to be generated.	`detail`

Table 75: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Errors in period	Number of symbol errors in the period reported in the received event PDU.	detail
Total errors	Number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors.	detail
OAM Received Frame Error Event Information		
Events	Number of errored frame event TLVs that have been received after the OAM sublayer was reset.	detail
Window	Duration of the window in terms of the number of 100 ms period intervals.	detail
Threshold	Number of detected errored frames required for the event to be generated.	detail
Errors in period	Number of detected errored frames in the period.	detail
Total errors	Number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	detail
OAM Received Frame Period Error Event Information		
Events	Number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.	detail
Window	Duration of the frame seconds window.	detail
Threshold	Number of frame seconds errors in the period.	detail

Table 75: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Errors in period	Number of frame seconds errors in the period.	detail
Total errors	Number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.	detail
OAM Transmitted Symbol Error Event Information		
Events	Number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.	detail
Window	The symbol error event window in the transmitted PDU.	detail
Threshold	Number of errored symbols in the period required for the event to be generated.	detail
Errors in period	Number of symbol errors in the period reported in the transmitted event PDU.	detail
Total errors	Number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.	detail
OAM Transmitted Frame Error Event Information		
Events	Number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.	detail
Window	Duration of the window in terms of the number of 100-ms period intervals.	detail
Threshold	Number of detected errored frames required for the event to be generated.	detail
Errors in period	Number of detected errored frames in the period.	detail

Table 75: show oam ethernet link-fault-management Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Total errors	Number of errored frames that have been detected after the OAM sublayer was reset.	detail

Sample Output

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
  Status: Running, Discovery state: Send Any
  Peer address: 2001:bd8:00:31
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
  Remote loopback status: Disabled on local port, Enabled on peer port
  Remote entity information:
    Remote MUX action: discarding, Remote parser action: loopback
    Discovery mode: active, Unidirectional mode: unsupported
    Remote loopback mode: supported, Link events: supported
    Variable requests: unsupported
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
  Interface: ge-0/0/1
    Status: Running, Discovery state: Send Any
    Peer address: 2001:bd8:00:31
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 186365, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM transmit statistics:
      Information: 186347, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM received symbol error event information:
```

```
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame period error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM transmitted symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM transmitted frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: supported, Link events: supported
      Variable requests: unsupported
```

Release Information

Statement for SRX Series devices introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

clear oam ethernet connectivity-fault-management path-database | 775

clear oam ethernet connectivity-fault-management statistics

Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways

Example: Configuring Ethernet OAM Link Fault Management on a Security Device

show poe controller (View)

IN THIS SECTION

- Syntax | 1080
- Description | 1080
- Options | 1080
- Required Privilege Level | 1080
- Output Fields | 1081
- Sample Output | 1081
- Release Information | 1081

Syntax

```
show poe controller
```

Description

Display the status of the Power over Ethernet (PoE) controller.

Options

none—Display general parameters of the PoE software module controller.

Required Privilege Level

View

Output Fields

Table 76 on page 1081 lists the output fields for the `show poe controller` command. Output fields are listed in the approximate order in which they appear.

Table 76: show poe controller Output Fields

Field name	Field Description
Controller-index	Identifies the controller.
Maximum-power	Specifies the maximum power that can be provided by the SRX Series device to PoE ports.
Power-consumption	Specifies the total amount of power allocated to the PoE ports.
Guard-band	Shows the guard band configured on the controller.
Management	Shows the power management mode.

Sample Output

show poe controller

`user@host>`**show poe controller**

```
Controller  Maximum   Power              Guard band  Management
index       power     consumption
   0        150.0 W    0.0 W               0 W         Static
```

Release Information

Command introduced in Junos OS Release 9.5.

show pppoe interfaces

IN THIS SECTION

- Syntax | 1082
- Description | 1082
- Options | 1082
- Required Privilege Level | 1083
- Output Fields | 1083
- Sample Output | 1086
- Release Information | 1088

Syntax

```
show pppoe interfaces
<brief | detail | extensive>
<pp0.logical>
```

Description

Display session-specific information about PPPoE interfaces.

Options

none	Display interface information for all PPPoE interfaces.
brief | detail	(Optional) Display the specified level of output.

extensive (Optional) Display information about the number of packets sent and received and the number of timeouts during a PPPoE session.

pp0.*logical* (Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16,385. The logical unit number for dynamic interfaces can be a value from 1,073,741,824 through the maximum number of logical interfaces supported on your SRX300, SRX320, and SRX340, and SRX550M devices.

Required Privilege Level

view

Output Fields

Table 77 on page 1083 lists the output fields for the `show pppoe interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 77: show pppoe interfaces Output Fields

Field Name	Field Description
`Index`	Index number of the logical interface, which reflects its initialization sequence.
`State`	State of the logical interface: `up` or `down`.
`Session ID`	Session ID.
`Service name`	Type of service required (can be used to indicate an ISP name, a class, or quality of service).
`Configured AC name`	Configured access concentrator name.
`Session AC name`	Name of the access concentrator.

Table 77: show pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description
`Remote MAC address` or `Remote MAC`	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.
`Auto-reconnect timeout`	Timeout value for reconnecting after a PPPoE session is terminated (in seconds).
`Idle timeout`	Length of time (in seconds) that a connection can be idle before disconnecting.
`Session uptime`	Length of time the session has been up, in *hh:mm:ss*.
`Ignore End-Of-List tag`	Disables the `End-of-List` tag to continue processing of other tags after the `End-of-List` tag in a PPPoE Active Discovery Offer (PADO) packet.
`Underlying interface`	Interface on which PPPoE is running.

Table 77: show pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description
`Packet Type`	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • `PADI`—PPPoE Active Discovery Initiation packets. • `PADO`—PPPoE Active Discovery Offer packets. • `PADR`—PPPoE Active Discovery Request packets. • `PADS`—PPPoE Active Discovery Session-Confirmation packets. • `PADT`—PPPoE Active Discovery Termination packets. • `Service name error`—Packets for which the Service-Name request could not be honored. • `AC system error`—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • `Generic error`—Packets that indicate an unrecoverable error occurred. • `Malformed packets`—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • `Unknown packets`—Unrecognized packets.
`Timeout`	Timeouts that occur during the PPPoE session: • `PADI`—No PADI packets received within the timeout period. • `PADO`—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • `PADR`—No PADR packets received within the timeout period.

Table 77: show pppoe interfaces Output Fields *(Continued)*

Field Name	Field Description
Receive Error Counters	Error counters received during the PPPoE session: • PADI—No PADI error counters received during the session. • PADO—No PADO error counters received during the session. • PADR—No PADR error counters received during the session. • PADS—No PADS error counters received during the session.

Sample Output

show pppoe interfaces

```
user@host> show pppoe interfaces
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
```

show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
Interface       Underlying           State       Session    Remote
                interface                        ID         MAC
pp0.0           ge-0/0/1.0           Session up  4          b0:c6:9a:74:5e:c1
```

show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  Ignore End-Of-List tag: Enable
```

show pppoe interfaces extensive

```
user@host> show pppoe interfaces extensive
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:22 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  PacketType                       Sent         Received
    PADI                              1                0
    PADO                              0                1
    PADR                              1                0
    PADS                              0                1
    PADT                              0                0
    Service name error                0                0
    AC system error                   0                0
    Generic error                     0                0
    Malformed packets                 0                0
    Unknown packets                   0                0
  Timeout
    PADI                              0
    PADO                              0
    PADR                              0
  Receive Error Counters
```

```
    PADI                        0
    PADO                        0
    PADR                        0
    PADS                        0
```

Release Information

Command introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

Ethernet Interfaces Overview | 160

show pppoe statistics

IN THIS SECTION

- Syntax | 1089
- Description | 1089
- Options | 1089
- Required Privilege Level | 1089
- Output Fields | 1089
- Sample Output | 1091
- Release Information | 1092

Syntax

```
show pppoe statistics
<logical-interface-name>
```

Description

Display statistics information about PPPoE interfaces.

Options

none	Display PPPoE statistics for all interfaces.
logical-interface-name	(Optional) Name of an underlying PPPoE logical interface.

Required Privilege Level

view

Output Fields

Table 78 on page 1089 lists the output fields for the `show pppoe statistics` command. Output fields are listed in the approximate order in which they appear.

Table 78: show pppoe statistics Output Fields

Field Name	Field Description
`Active PPPoE sessions`	Total number of active PPPoE sessions.

Table 78: show pppoe statistics Output Fields *(Continued)*

Field Name	Field Description
`Packet Type`	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: • `PADI`—PPPoE Active Discovery Initiation packets. • `PADO`—PPPoE Active Discovery Offer packets. • `PADR`—PPPoE Active Discovery Request packets. • `PADS`—PPPoE Active Discovery Session-Confirmation packets. • `PADT`—PPPoE Active Discovery Termination packets. • `Service name error`—Packets for which the Service-Name request could not be honored. • `AC system error`—Packets for which the access concentrator experienced an error in performing the host request. For example, the host had insufficient resources to create a virtual circuit. • `Generic error`—Packets that indicate an unrecoverable error occurred. • `Malformed packets`—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • `Unknown packets`—Unrecognized packets.
`Timeout`	Timeouts that occur during the PPPoE session: • `PADI`—No PADI packets received within the timeout period. • `PADO`—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • `PADR`—No PADR packets received within the timeout period.

Table 78: show pppoe statistics Output Fields *(Continued)*

Field Name	Field Description
Receive Error Counters	Error counters received during the PPPoE session: • PADI—No PADI error counters received during the session. • PADO—No PADO error counters received during the session. • PADR—No PADR error counters received during the session. • PADS—No PADS error counters received during the session.

Sample Output

show pppoe statistics

```
user@host> show pppoe statistics
Active PPPoE sessions: 0

  PacketType                      Sent         Received
    PADI                             0                0
    PADO                             0                0
    PADR                             0                0
    PADS                             0                0
    PADT                             0                0
    Service name error               0                0
    AC system error                  0                0
    Generic error                    0                0
    Malformed packets                0                0
    Unknown packets                  0                0
  Timeout
    PADI                             0
    PADO                             0
    PADR                             0
  Receive Error Counters
    PADI                             0
    PADO                             0
```

```
    PADR                    0
    PADS                    0
```

Release Information

Command is t introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

- show pppoe interfaces | 1082
- Ethernet Interfaces Overview | 160

show poe telemetries

IN THIS SECTION

- Syntax | 1093
- Description | 1093
- Options | 1093
- Required Privilege Level | 1093
- Output Fields | 1093
- Sample Output | 1094
- Release Information | 1095

Syntax

```
show poe telemetries
<interface interface-name count number>
<count number interface interface-name>
```

Description

Display a history of power consumption on the specified interface. Telemetries must be enabled on the interface before you can display a history of power consumption.

Options

- `Interface` *interface-name*—Display telemetries for the specified PoE interface.
- `count` *number*—Display the specified number of telemetries records for the specified PoE interface.

Required Privilege Level

View

Output Fields

Table 79 on page 1093 lists the output fields for the `show poe telemetries interface` command. Output fields are listed in the approximate order in which they appear.

Table 79: show poe telemetries interface Output Fields

Field name	Field Description
S1 No	Number of the record for the specified port. The last record is the most is the most recent.

Table 79: show poe telemetries interface Output Fields *(Continued)*

Field name	Field Description
Timestamp	Time that the power-consumption data was gathered.
Power	Amount of power provided by the specified port at the time the data was gathered.
Voltage	Voltage on the specified port at the time the data was gathered.

Sample Output

show poe telemetries interface

user@host>**show poe telemetries interface ge-0/0/1 count 8**

```
Sl No    Timestamp                Power    Voltage
   1     Fri Jan 04 11:41:15 2009 6.6 W    47.2 V
   2     Fri Jan 04 11:40:15 2009 6.6 W    47.2 V
   3     Fri Jan 04 11:39:15 2009 6.6 W    47.2 V
   4     Fri Jan 04 11:38:15 2009 6.6 W    47.2 V
   5     Fri Jan 04 11:37:15 2009 6.6 W    47.2 V
   6     Fri Jan 04 11:36:15 2009 6.6 W    47.2 V
   7     Fri Jan 04 11:35:15 2009 6.6 W    47.2 V
   8     Fri Jan 04 11:34:15 2009 6.6 W    47.2 V
```

user@host>**show poe telemetries count 5 interface ge-0/0/1**

```
Sl No    Timestamp                Power    Voltage
   1     Fri Jan 04 11:47:15 2009 6.6 W    47.2 V
   2     Fri Jan 04 11:38:15 2009 6.6 W    47.2 V
   3     Fri Jan 04 11:29:15 2009 6.6 W    47.2 V
   4     Fri Jan 04 11:11:15 2009 6.6 W    47.2 V
```

```
    5     Fri Jan 04 11:10:15 2009 6.6 W    47.2 V
```

Release Information

Command modified in Junos OS Release 12.3X48-D10.

show services accounting

IN THIS SECTION

- Syntax | **1095**
- Description | **1096**
- Options | **1096**
- Required Privilege Level | **1097**
- Output Fields | **1097**
- Sample Output | **1097**
- Release Information | **1098**

Syntax

```
show services accounting
aggregation
errors
<inline-jflow | inline-jflow fpc-slot slot number>
flow
<inline-jflow | inline-jflow fpc-slot slot number>
flow-detail
memory
packet-size-distribution
```

```
status
<inline-jflow | inline-jflow fpc-slot slot number>
usage
```

Description

Display sampled accounting service.

Options

- aggregation—Display aggregation information.
- errors —Display error statistics.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- flow—Display flow information.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- flow-detail—Display flow detail.
- memory—Display memory information.
- packet-size-distribution—Display packet size distribution.
- status–Display service accounting parameters.
  - inline-jflow — Display service accounting inline flow monitoring parameters.
  - fpc-slot *slot number*— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring.
- usage–Display CPU usage.

Required Privilege Level

view

Output Fields

Lists the output fields for the **show services accounting** command.

Sample Output

show services accounting status inline-jflow

```
user@host> show services accounting status inline-jflow
Status information
    FPC Slot: 5
    Export format: IP-FIX(V9)
    IPv4 Route Record Count: 16, IPv6 Route Record Count: 5
    Route Record Count: 21, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes
```

show services accounting errors inline-jflow

```
user@host> show services accounting errors inline-jflow
Error Information
  FPC Slot:  5
  PIC Slot:  0
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0
  AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No

  IPv4 Errors:
  IPv4 Flow Creation Failures: 0
  IPv4 Route Record Lookup Failures: 0
```

```
  IPv4 AS Lookup Failures: 0
  IPv4 Export Packet Failures: 0

  IPv6 Errors:
  IPv6 Flow Creation Failures: 0
  IPv6 Route Record Lookup Failures: 0
  IPv6 AS Lookup Failures: 0
  IPv6 Export Packet Failures: 0
```

show service accounting flow inline-jflow

```
user@host> show service accounting flow inline-jflow
Flow Information
FPC Slot:  5
PIC Slot:  0
Flow Packets: 2  Flow Bytes: 0
Active Flows: 1  Total Flows: 2
Flows Exported: 0 Flow Packets Exported: 231
Flows Inactive Timed Out: 1 Flows Active Timed Out: 2

IPv4 Flows:
IPv4 Flow Packets: 1  IPv4 Flow Bytes: 0
IPv4 Active Flows: 1  IPv4 Total Flows: 1
IPv4 Flows Exported: 0 IPv4 Flow Packets Exported: 132
IPv4 Flows Inactive Timed Out: 0 IPv4 Flows Active Timed Out: 1

IPv6 Flows:
IPv6 Flow Packets: 1  IPv6 Flow Bytes: 0
IPv6 Active Flows: 0  IPv6 Total Flows: 1
IPv6 Flows Exported: 0 IPv6 Flow Packets Exported: 99
IPv6 Flows Inactive Timed Out: 1 IPv6 Flows Active Timed Out: 1
```

Release Information

Command introduced in Junos OS Release 10.4. The **inline-jflow** and **fpc-slot** options are added in Junos OS Release 12.1X45-D10.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

show services accounting aggregation (View)

IN THIS SECTION

- Syntax | 1099
- Description | 1099
- Options | 1099
- Required Privilege Level | 1100
- Release Information | 1100

Syntax

```
show services accounting aggregation
```

Description

Display aggregation information for the accounting service.

Options

- as—Display aggregation type AS.
- destination-prefix—Display aggregation type destination-prefix.
- protocol-port—Display aggregation type protocol-port.

- source-destination-prefix—Display aggregation type source-destination-prefix.
- source-prefix—Display aggregation type source-prefix.
- template—Display aggregation type template.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

show services accounting aggregation template (View)

IN THIS SECTION

- Syntax | **1101**
- Description | **1101**
- Options | **1101**
- Required Privilege Level | **1101**
- Release Information | **1101**

Syntax

```
show services accounting aggregation template
```

Description

Display aggregation type template.

Options

- detail—Display detailed output.
- extensive—Display extensive output.
- template-name—Display name of the template.
- terse—Display terse output (default).

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

show services accounting flow-detail (View)

IN THIS SECTION

- Syntax | 1102
- Description | 1102
- Options | 1102
- Required Privilege Level | 1103
- Release Information | 1103

Syntax

```
show services accounting flow-detail
```

Description

Display flow detail

Options

- destination-as—Filter term destination AS.
- destination-port—Filter term destination port.
- destination-prefix—Filter term destination prefix.
- detail—Display detailed output.
- extensive–Display extensive output.
- input-snmp-interface-index–Filter term input SNMP interface index.

- limit–Display maximum number of flows to display.
- name-Display name of the service, wildcard, or “all”.
- order–Display order for displaying flows.
- output-snmp-interface-index–Filter term output SNMP interface index.
- proto–Filter term protocol.
- source-as–Filter term source AS.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Flow Aggregation to Use Version 9 Flow Templates

show wlan access-points

IN THIS SECTION

- Syntax | **1104**
- Description | **1104**
- Options | **1104**
- Required Privilege Level | **1105**
- Output Fields | **1105**

Sample Output | 1108
Release Information | 1111

Syntax

Syntax (SRX Series devices)

```
show wlan
<detail>
<virtual-access-points>
<client-associations>
<neighbors>
<summary>
<radio>
<ha>
```

Description

Display information about wireless and virtual access point WLANs configured on the wireless LAN interface wl-*x*/0/0.

Options

detail	(Optional) Display the specified level of output.
virtual access points	Display Virtual access Points (VAPs) status and statistics.
client-associations	Client association number of the specified access point.
neighbors	List neighboring access point information including the MAC address, WPA, band, channel and SSID values.

summary Display the access point configuration summary as output.

radio Display access point radio information.

ha Display wireless access point status in chassis cluster mode. In chassis cluster mode, only one wirelesss interface is active and the wireless client can choose the required secure method to establish the wireless connection. After the wireless interface failover, wireless client retries the secure process to establish the new wireless connection.

In WIFI mPIM HA mode, there is one WAP mPIM card on each node. The WAP mPIM cards on both nodes have the same WAP configuration such as SSID, channel, and bandwidth. WAP0 is active on node0 and WAP1 is inactive on node1. Users are connected to WAP0 initially. When there is WAP0 failure, WAP1 is changed to active and the users are connected to WAP1 automatically.

The `show wlan access-points ha status` command is supported on primary routing engine. This command displays information from WAP mPIM card on both nodes.

Required Privilege Level

view

Output Fields

Table 80 on page 1105 lists the output fields for the `show wlan access-points` command. Output fields are listed in the approximate order in which they appear.

Table 80: `show wlan access-points` output fields

Field Name	Field Description	Level of Output
Access point	Name of the wireless access point.	All levels.
Type	Internal.	`detail`

Table 80: `show wlan access-points` output fields *(Continued)*

Field Name	Field Description	Level of Output
Location	Location of the access point.	`detail`
Serial Number	Serial number of the Mini-PIM.	`detail`
Firmware Version	Firmware version of the Mini-PIM.	`detail`
Alternate Version	Alternate firmware version of the Mini-PIM.	`detail`
Country	Country code.	`detail`
Access Interface	Name of the wireless LAN interface on the WI-Fi Mini-PIM.	`detail`
Packet Capture	Status of the traffic flow.	`detail`
Capture Interface	Interface captured on radio.	`detail`
Capture File	Capture file name.	`detail`
Capture Duration	Duration of capture in seconds.	`detail`
Capture File Size	File size of capture in kilobytes.	`detail`
MAC Address	MAC address of the Ethernet port.	`detail`
IPv4 Address	IPv4 address of the access point.	`detail`
Mode	Authentication mode on the radio.	`detail`

Table 80: `show wlan access-points` output fields *(Continued)*

Field Name	Field Description	Level of Output
Channel	Channel bandwidth on the radio.	`detail` `neighbors`
VAP	Name of the virtual access point.	
SSID	Network name of the virtual access point.	`detail` `neighbors`
VLAN ID	VLAN ID associated with the access point.	`detail`
Traffic Statistics	Number and rate of bytes and packets received and transmitted on the wireless LAN interface. • Input bytes—Number of bytes received on the wireless LAN interface. • Output bytes—Number of bytes transmitted on the wireless LAN interface. • Input packets—Number of packets received on the wireless LAN interface. • Output packets—Number of packets transmitted on the wireless LAN interface.	`detail`
Client number	Client number on Radios 1 and 2.	`client-associations`
MAC Privacy	Client MAC address as per the virtual access.	`neighbors`
WPA	Status of the security authentication method.	`neighbors`
Band	Status of the band.	`neighbors`

Sample Output

show wlan access-points (SRX Series devices)

```
user@router> show wlan access-points
Active access points information
Access-Point     Type   Interface    Radio-mode/Channel/Bandwidth
bj340d-ha        Int    wl-2/0/0     Off/Off/Off, Off/Off/Off
```

show wlan access-points (SRX Series devices)

```
user@router> show wlan access-points e09-22-ha-ap
node0:
--------------------------------------------------------------------------

Active access point information

Access Point        : e09-22-ha-ap
Type                : Internal
Access Interface    : wl-2/0/0
Packet Capture      : Off
Radio1              : Mode: IEEE 802.11a/n/ac, Channel: 153, Bandwidth: 40
Radio2              : Mode: Off, Channel: Off, Bandwidth: Off

node1:
--------------------------------------------------------------------------

Active access point information

Access Point        : e09-22-ha-ap
Type                : Internal
Access Interface    : wl-7/0/0
Packet Capture      : Off
Radio1              : Mode: IEEE 802.11a/n/ac, Channel: 153, Bandwidth: 40
Radio2              : Mode: Off, Channel: Off, Bandwidth: Off
```

show wlan access-points detail (SRX Series devices)

```
user@router> show wlan access-points bj340d-ha detail
Active access point detail information
Access Point          : bj340d-ha
Type                  : Internal
Location              : Default Location
Serial Number         : EV0519AF0022
Firmware Version      : v1.2.2
Alternate Version     : v1.1.8
Country               : US
Access Interface      : wl-2/0/0
System Time           : Thu Oct 31 12:11:40 UTC 2019
Packet Capture        : Off
Ethernet Port:
MAC Address           : 94:f7:ad:2c:7b:87

Radio1:
                        Status            : Off
                        MAC Address       : 94:f7:ad:2c:7b:89
                        Temperature       : 0
                        Mode              : Off
                        Channel           : Off
                        Bandwidth         : Off
                        Transmit Power    : Off

Radio2:
                        Status            : Off
                        MAC Address       : 94:f7:ad:2c:7b:88
                        Temperature       : 0
                        Mode              : Off
                        Channel           : Off
                        Bandwidth         : Off
                        Transmit Power    : Off
```

show wlan access-points virtual access-points (SRX Series devices)

```
user@host> show wlan access-points bj340d-ha virtual-access-points
Virtual access points information
Access Point           : bj340d-ha
```

```
Radio1:
                    VAP0:
                    SSID                : juniper_ap_0
                    MAC Address         : 00:11:22:33:44:55
                    VLAN ID             : 5
                    Traffic Statistics  :
                    Input Bytes         : 37979930
                    Output Bytes        : 54231321
                    Input Packets       : 210737
                    Output Packets      : 298451

                    VAP1:
                    SSID                : juniper_ap_3
                    MAC Address         : 00:11:22:33:44:58
                    VLAN ID             : 15
                    Traffic Statistics  :
                    Input Bytes         : 83421130
                    Output Bytes        : 21221311
                    Input Packets       : 510837
                    Output Packets      : 238451

Radio2:
                    VAP0:
                    SSID                : juniper_ap_7
                    MAC Address         : 00:11:22:33:44:75
                    VLAN ID             : 5
                    Traffic Statistics  :
                    Input Bytes         : 67079335
                    Output Bytes        : 34932321
                    Input Packets       : 415737
                    Output Packets      : 308451

                    VAP1:
                    SSID                : juniper_ap_3
                    MAC Address         : 00:11:22:33:44:78
                    VLAN ID             : 49
                    Traffic Statistics  :
                    Input Bytes         : 37979930
                    Output Bytes        : 54231321
                    Input Packets       : 512837
                    Output Packets      : 901151
```

show wlan access-points client-associations summary (SRX Series Devices)

```
user@host> show wlan access-points bj340d-ha client associations summary
Access point client associations summary
Access point                        : bj340d-ha

Client number on radio 1 (5.0 GHz)  : 30
Client number on radio 2 (2.4 GHz)  : 20
Total client number on access point : 50
```

show wlan access-points neighbors (SRX Series Devices)

```
user@host> show wlan access-points bj340d-ha neighbors
Access point neighbors information
Access point         : bj340d-ha

MAC                 Privacy   WPA    Band    Channel     SSID
00:11:22:33:44:55   Off       Off    2.4     10          xyz-ap
00:12:23:45:56:67   On        On     5       100         abc-ap
```

show wlan access-points ha status (SRX Series Devices)

```
user@host> show wlan access-points ha status

Acess-point                              Interface           HA-status

e09-22-ha-ap                           wl-2/0/0               Inactive
e09-22-ha-ap                           wl-7/0/0               Active
```

Release Information

Command introduced in Junos OS Release 19.4R1.

ha option is introduced in the Junos OS Release 20.3R1 for SRX Series devices.

RELATED DOCUMENTATION

Wi-Fi Mini-Physical Interface Module Overview | 443

Configure Wi-Fi Mini-PIM | 446

wlan | 767

speed (Chassis Cluster)

IN THIS SECTION

- Syntax | 1112
- Description | 1112
- Options | 1113
- Required Privilege Level | 1113
- Release Information | 1114

Syntax

```
set chassis cluster control-port speed (1g |10g);
```

Description

The SRX4600 supports three different PIC types—8-port 10-Gigabit Ethernet PIC, 4-port 40-Gigabit or 100-Gigabit Ethernet PIC, and 4-port 10-Gigabit Ethernet PIC (in a chassis cluster). Out of the four ports on the 10-Gigabit Ethernet PIC in a chassis cluster, two ports are fabric ports and the other two ports are chassis cluster control ports. The two fabric ports do not support 1-Gbps speed. Only the two control ports of the chassis cluster support a port speed of 1 Gbps. If you replace one node with return merchandise authorization (RMA), then you should set the speed and reboot the node for the control-link to come online.

On chassis cluster control interfaces, you can configure the operating speed of the 4-port 10-Gigabit Ethernet PIC from default 10-Gbps port speed to 1-Gbps port speed. You must reboot the device for the changed configuration to take effect.

The chassis cluster control interfaces do not support multiple speeds.

Following are the list of optics supported on SRX4600:

- SRX-SFP-1GE-LX
- SRX-SFP-1GE-LX-ET
- SRX-SFP-1GE-SX
- SRX-SFP-1GE-SX-ET

Autonegotiation is automatically disabled when 1-Gbps speed is configured on the interfaces.

NOTE:

- The interface name for any xe interface remains same after converting its speed from 10G to 1G.
- To view the speed configured for the interface, execute the `show interfaces extensive` command. The `Speed Configuration` field's value of `1G` or `AUTO` in the command output indicates whether the current operation speed of the interface is 1 Gbps or the default 10 Gbps, respectively.

Options

- 1g — Link speed of 1 Gbps
- 10g — Link speed of 10 Gbps

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1 for SRX4600.

RELATED DOCUMENTATION

speed (Gigabit Ethernet interface) | 740