



Junos[®] OS

Interfaces Feature Guide for Security Devices



Modified: 2018-09-03

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. and/or its affiliates in the United States and other countries. All other trademarks may be 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 Feature Guide for Security Devices
Copyright © 2018 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://www.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xxiii
	Documentation and Release Notes	xxiii
	Using the Examples in This Manual	xxiii
	Merging a Full Example	xxiv
	Merging a Snippet	xxiv
	Documentation Conventions	xxv
	Documentation Feedback	xxvii
	Requesting Technical Support	xxvii
	Self-Help Online Tools and Resources	xxviii
	Opening a Case with JTAC	xxviii
Part 1	Overview	
Chapter 1	Introduction to Interfaces	3
	Understanding Interfaces	3
	Network Interfaces	4
	Services Interfaces	5
	Special Interfaces	8
	Interface Naming Conventions	8
	Understanding the Data Link Layer	10
	Physical Addressing	11
	Network Topology	11
	Error Notification	11
	Frame Sequencing	11
	Flow Control	11
	Data Link Sublayers	11
	MAC Addressing	12
	Configuring IOC to NPC Mapping	12
	Monitoring Interfaces	13
	Understanding GRE Keepalive Time	15
	Configuring GRE Keepalive Time	15
	Configuring Keepalive Time and Hold time for a GRE Tunnel Interface	16
	Display GRE Keepalive Time Configuration	16
	Display Keepalive Time Information on a GRE Tunnel Interface	17
	SRX4600 LED Scheme Overview	19
Chapter 2	Configuring Interface Logical Properties	21
	Understanding Interface Logical Properties	21
	Understanding Protocol Families	22
	Common Protocol Suites	22
	Other Protocol Suites	22

	Understanding IPv4 Addressing	23
	IPv4 Classful Addressing	23
	IPv4 Dotted Decimal Notation	24
	IPv4 Subnetting	24
	IPv4 Variable-Length Subnet Masks	25
	Understanding IPv6 Address Space, Addressing, Address Format, and Address Types	26
	Understanding IP Version 6 (IPv6)	26
	Understanding IPv6 Address Types and How Junos OS for SRX Series Services Gateway Uses Them	26
	IPv6 Address Scope	27
	IPv6 Address Structure	27
	Understanding IPv6 Address Space, Addressing, and Address Types	28
	Understanding IPv6 Address Format	28
	Limitations	29
	Configuring the inet6 IPv6 Protocol Family	29
	Enabling Flow-Based Processing for IPv6 Traffic	30
	Configuring Flow Aggregation to Use Version 9 Flow Templates	32
	Configuring the Traffic to Be Sampled	32
	Configuring the Version 9 Template Properties	33
	Restrictions	34
	Fields Included in Each Template Type	35
	inet Sampling Behavior	37
	Verification	37
	Examples: Configuring Version 9 Flow Templates	38
	Understanding IPv6 Support VDSL2 Interfaces	41
	Example: Configuring the IPv6 Address on an ADSL Interface	41
	Understanding MAC Limiting on Layer 3 Routing Interfaces	44
	Overview	44
	Limitations	46
Chapter 3	Understanding Interface Physical Properties	47
	Understanding Interface Physical Properties	47
	Understanding Bit Error Rate Testing	48
	Understanding Interface Clocking	49
	Data Stream Clocking	50
	Explicit Clocking Signal Transmission	50
	Understanding Frame Check Sequences	50
	Cyclic Redundancy Checks and Checksums	50
	Two-Dimensional Parity	51
	MTU Default and Maximum Values	51
	Understanding Jumbo Frames Support for Ethernet Interfaces	53
Chapter 4	Configuring VLAN Tagging	55
	Understanding Virtual LANs	55
	VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices	57
	Configuring VLAN Tagging	57
	Configuring Single-Tag Framing	58
	Configuring Dual Tagging	58

	Configuring Mixed Tagging	58
	Configuring Mixed Tagging Support for Untagged Packets	59
Chapter 5	Rate-Selectability	61
	SRX4600 Gateway Rate-Selectability Overview	61
	User-Configurable Rate Selectability of SRX4600 Services Gateway	64
	Maximum number of 10/40/100GE ports Configurable at PIC and Port Mode	65
	Port Configuration - PIC Level	65
	Configuring 40-Gigabit Ethernet ports to 4X10-Gigabit Ethernet using Breakout Cables	66
	Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on SRX4600 Gateway	67
	Invalid Port Configuration	68
	Configuring Active Ports on SRX4600 Gateway with Rate Selectability	68
Part 2	Configuring DS1 Interfaces	
Chapter 6	Configuring DS1 Interfaces	73
	Understanding T1 and E1 Interfaces	73
	T1 Overview	73
	E1 Overview	74
	T1 and E1 Signals	74
	Encoding	74
	AMI Encoding	74
	B8ZS and HDB3 Encoding	75
	T1 and E1 Framing	75
	ESF Framing for T1	75
	T1 and E1 Loopback Signals	75
	Example: Configuring a T1 Interface	76
	Example: Deleting a T1 Interface	79
Chapter 7	Configuring DS3 Interfaces	81
	Understanding T3 and E3 Interfaces	81
	Multiplexing DS1 Signals	81
	DS2 Bit Stuffing	82
	DS3 Framing	82
	M13 Asynchronous Framing	82
	C-Bit Parity Framing	84
	Example: Configuring a T3 Interface	86
	Example: Deleting a T3 Interface	88
Chapter 8	Configuring DS3 Interfaces	91
	Understanding T3 and E3 Interfaces	91
	Multiplexing DS1 Signals	91
	DS2 Bit Stuffing	92

	DS3 Framing	92
	M13 Asynchronous Framing	92
	C-Bit Parity Framing	94
	Example: Configuring a T3 Interface	96
	Example: Deleting a T3 Interface	98
Chapter 9	Configuring 1-Port Clear Channel DS3/E3 GPIM	101
	Understanding the 1-Port Clear Channel DS3/E3 GPIM	101
	Supported Features	101
	Interface Naming	102
	Physical Interface Settings	102
	Logical Interface Settings	102
	Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode	104
	Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode	106
	Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode	107
Part 3	Configuring DSL Interfaces	
Chapter 10	Configuring ADSL Interfaces	111
	ADSL Interface Overview	111
	ADSL Systems	112
	ADSL2 and ADSL2+	113
	ATM CoS Support	113
	ADSL and SHDSL Interfaces Configuration Overview	114
	Example: Configuring ATM-over-SHDSL Network Interfaces	118
	Example: Configuring MLPPP-over-ADSL Interfaces	124
	Example: Configuring the DHCP Client on ADSL Interface	126
	Example: Configuring CHAP on DSL Interfaces	130
	Example: Configuring ATM-over-ADSL Network Interfaces	138
Chapter 11	Configuring G.SHDSL Interfaces	145
	SHDSL Interface Overview	145
	G.SHDSL Mini-PIM Overview	146
	Operating Modes and Line Rates of the G.SHDSL Mini-PIM	148
	G.SHDSL Mini-PIM Configuration Overview	148
	Example: Configuring the G.SHDSL Interface	150
	Example: Configuring the G.SHDSL Interface on SRX Series Devices	158
	Example: Configuring the G.SHDSL Interface in EFM Mode	170
Chapter 12	Configuring VDSL2 Interfaces	181
	VDSL2 Interface Technology Overview	181
	VDSL2 Vectoring Overview	182
	VDSL2 Network Deployment Topology	182
	VDSL2 Interface Support on SRX Series Devices	184
	VDSL2 Interface Compatibility with ADSL Interfaces	184
	VDSL2 Interfaces Supported Profiles	186

	VDSL2 Interfaces Supported Features	187
	Example: Configuring VDSL2 Interfaces in ADSL Mode (Basic)	187
	Example: Configuring VDSL2 Interfaces in ADSL Mode (Detail)	193
	Example: Configuring VDSL2 Interfaces (Basic)	220
	Example: Configuring VDSL2 Interfaces (Detail)	227
	Upgrading the VDSL PIC Firmware	253
Part 4	Configuring Ethernet Interfaces	
Chapter 13	Performing Initial Configuration on Ethernet Interfaces	259
	Understanding Ethernet Interfaces	259
	Ethernet Access Control and Transmission	260
	Collisions and Detection	260
	Collision Detection	260
	Backoff Algorithm	260
	Collision Domains and LAN Segments	261
	Repeaters	261
	Bridges and Switches	261
	Broadcast Domains	262
	Ethernet Frames	262
	Understanding Static ARP Entries on Ethernet Interfaces	263
	Understanding Promiscuous Mode on Ethernet Interface	263
	Understanding Promiscuous Mode on the SRX5K-MPC	264
	Understanding Port Mirroring on SRX Devices	264
	Example: Creating an Ethernet Interface	265
	Example: Deleting an Ethernet Interface	266
	Example: Configuring Static ARP Entries on Ethernet Interfaces	267
	Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)	270
	Example: Configuring Promiscuous Mode on the SRX5K-MPC	271
	Configuring Port Mirroring on SRX Devices	275
Chapter 14	Configuring Aggregated Ethernet Interfaces	279
	Understanding Aggregated Ethernet Interfaces	279
	LAGs	280
	LACP	280
	Aggregated Ethernet Interfaces Configuration Overview	282
	Understanding the Aggregated Ethernet Interfaces Device Count	282
	Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device	283
	Understanding Physical Interfaces for Aggregated Ethernet Interfaces	284
	Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces	285
	Understanding Aggregated Ethernet Interface Link Speed	286
	Example: Configuring Aggregated Ethernet Link Speed	286
	Understanding Minimum Links for Aggregated Ethernet Interfaces	287
	Example: Configuring Aggregated Ethernet Minimum Links	288
	Understanding Aggregated Ethernet Interface Removal	289
	Example: Deleting Aggregated Ethernet Interfaces	289

	Example: Deleting Aggregated Ethernet Interface Contents	290
	Verifying Aggregated Ethernet Interfaces	291
	Verifying Aggregated Ethernet Interfaces (terse)	291
	Verifying Aggregated Ethernet Interfaces (extensive)	292
	Understanding VLAN Tagging for Aggregated Ethernet Interfaces	293
	Understanding Promiscuous Mode for Aggregated Ethernet Interfaces	293
Chapter 15	Configuring Link Aggregation Control Protocol	295
	Understanding LACP on Standalone Devices	295
	Example: Configuring Link Aggregation Control Protocol (CLI Procedure)	296
	Verifying LACP on Standalone Devices	300
	Verifying LACP Statistics	300
	Verifying LACP Aggregated Ethernet Interfaces	300
	Understanding LACP on Chassis Clusters	301
	Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups	302
	Minimum Links	302
	Sub-LAGs	302
	Supporting Hitless Failover	303
	Managing Link Aggregation Control PDUs	303
	Example: Configuring LACP on Chassis Clusters	304
	Verifying LACP on Redundant Ethernet Interfaces	306
	LAG and LACP Support on SRX5000 Line Devices with I/O Cards (IOCs)	307
	LAG and LACP Support on the SRX5000 Module Port Concentrator	307
	LAG and LACP Support on the SRX5000 Line IOCs in Express Path Mode	308
	Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3	309
	Example: Configuring Aggregated Ethernet Device with LAG and LACP on a Security Device (CLI Procedure)	313
Chapter 16	Configuring Gigabit Ethernet Physical Interface Modules	317
	Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM	317
	Supported Features	317
	Interface Names and Settings	318
	Available Link Speeds and Modes	318
	Link Settings	319
	Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface	319
	Understanding the 2-Port 10-Gigabit Ethernet XPIM	325
	Supported Features	326
	Interface Names and Settings	327
	Copper and Fiber Operating Modes	327
	Link Speeds	327
	Link Settings	328
	Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface	328
	Understanding the 8-Port Gigabit Ethernet SFP XPIM	332
	Supported Features	333
	Interface Names and Settings	334
	Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs	334

Chapter 17	Configuring Port Mirroring	351
	Understanding Port Mirroring on SRX Devices	351
	Configuring Port Mirroring on SRX Devices	352
Chapter 18	Configuring Ethernet OAM Link Fault Management	355
	Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways	355
	Example: Configuring Ethernet OAM Link Fault Management on a Security Device	357
	Example: Configuring Remote Loopback Mode on VDSL Interfaces on a Security Device	361
Chapter 19	Configuring Power over Ethernet	367
	Understanding Power over Ethernet	367
	SRX Series Services Gateway PoE Specifications	367
	PoE Classes and Power Ratings	369
	PoE Options	370
	Example: Configuring PoE on All Interfaces	370
	Example: Configuring PoE on an Individual Interface	373
	Example: Disabling a PoE Interface	376
Part 5	Configuring Interface Encapsulation	
Chapter 20	Interface Encapsulation Overview	381
	Understanding Physical Encapsulation on an Interface	381
	Understanding Frame Relay Encapsulation on an Interface	382
	Virtual Circuits	382
	Switched and Permanent Virtual Circuits	382
	Data-Link Connection Identifiers	383
	Congestion Control and Discard Eligibility	383
	Understanding Point-to-Point Protocol	383
	Link Control Protocol	384
	PPP Authentication	384
	Network Control Protocols	385
	Magic Numbers	385
	CSU/DSU Devices	386
	Understanding High-Level Data Link Control	386
	HDLC Stations	386
	HDLC Operational Modes	387
Chapter 21	Configuring Point-to-Point Protocol over Ethernet	389
	Understanding Point-to-Point Protocol over Ethernet	389
	PPPoE Discovery Stage	390
	PPPoE Session Stage	391
	Understanding PPPoE Interfaces	392
	Example: Configuring PPPoE Interfaces	393
	Understanding PPPoE Ethernet Interfaces	399
	Example: Configuring PPPoE Encapsulation on an Ethernet Interface	400
	Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces	401
	Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface	402

	Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces	404
	Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface . .	404
	Understanding CHAP Authentication on a PPPoE Interface	407
	Example: Configuring CHAP Authentication on a PPPoE Interface	407
	Verifying Credit-Flow Control	409
	Verifying PPPoE Interfaces	410
	Verifying R2CP Interfaces	410
	Displaying Statistics for PPPoE	411
	Setting Tracing Options for PPPoE	412
Chapter 22	Configuring PPPoE-Based Radio-to-Router Protocol	415
	PPPoE-Based Radio-to-Router Protocols Overview	415
	Understanding the PPPoE-Based Radio-to-Router Protocol	416
	Configuring PPPoE-Based Radio-to-Router Protocols	418
	Example: Configuring the PPPoE-Based Radio-to-Router Protocol	418
	Credit Flow Control for PPPoE	421
	PPPoE Credit-Based Flow Control Configuration	421
Chapter 23	Configuring R2CP Radio-to-Router Protocol	423
	R2CP Radio-to-Router Protocol Overview	423
	Configuring the R2CP Radio-to-Router Protocol	424
Part 6	Configuring Link Services and Special Interfaces	
Chapter 24	Configuring Link Services Interfaces	431
	Link Services Interfaces Overview	431
	Services Available on a Link Services Interface	432
	Link Services Exceptions	432
	Configuring Multiclass MLPPP	433
	Queuing with LFI	434
	Queuing on Q2s of Constituent Links	434
	Compressed Real-Time Transport Protocol Overview	435
	Configuring Fragmentation by Forwarding Class	435
	Configuring Link-Layer Overhead	437
	Link Services Configuration Overview	437
	Verifying the Link Services Interface	438
	Verifying Link Services Interface Statistics	438
	Verifying Link Services CoS Configuration	441
	Troubleshooting the Link Services Interface	443
	Determine Which CoS Components Are Applied to the Constituent	
	Links	443
	Determine What Causes Jitter and Latency on the Multilink Bundle	445
	Determine If LFI and Load Balancing Are Working Correctly	445
	Determine Why Packets Are Dropped on a PVC Between a Juniper Networks	
	Device and a Third-Party Device	452
Chapter 25	Configuring Link Fragmentation and Interleaving	455
	Understanding Link Fragmentation and Interleaving Configuration	455
	Example: Configuring Link Fragmentation and Interleaving	456

Chapter 26	Configuring Class-of-Service on Link Services Interfaces	459
	Understanding How to Define Classifiers and Forwarding Classes	459
	Example: Defining Classifiers and Forwarding Classes	460
	Understanding How to Define and Apply Scheduler Maps	463
	Example: Configuring Scheduler Maps	465
	Understanding Interface Shaping Rates	468
	Example: Configuring Interface Shaping Rates	468
Chapter 27	Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles	471
	Understanding MLPPP Bundles and Link Fragmentation and Interleaving (LFI) on Serial Links	471
	Example: Configuring an MLPPP Bundle	472
Chapter 28	Configuring Multilink Frame Relay	477
	Understanding Multilink Frame Relay FRF.15	477
	Example: Configuring Multilink Frame Relay FRF.15	477
	Understanding Multilink Frame Relay FRF.16	480
	Example: Configuring Multilink Frame Relay FRF.16	481
Chapter 29	Configuring Compressed Real-Time Transport Protocol	485
	Understanding Compressed Real-Time Transport Protocol	485
	Example: Configuring the Compressed Real-Time Transport Protocol	485
Chapter 30	Configuring Link Services Queuing Interface	489
	Understanding the Internal Interface LSQ-0/0/0 Configuration	489
	Example: Upgrading from ls-0/0/0 to lsq-0/0/0 for Multilink Services	489
Chapter 31	Understanding Special Interfaces	493
	Understanding Management Interfaces	493
	Understanding the Discard Interface	494
	Understanding the Loopback Interface	494
	Configuring a Loopback Interface	495
Part 7	Configuring LTE Interfaces	
Chapter 32	Configuring LTE Interfaces	499
	LTE Mini-PIM Overview	499
	Supported Features	500
	Understanding the LTE Physical Interface	501
	Understanding the LTE Logical Interface	501
	Class of Service on the Dialer Interface	502
	LTE Mini-PIM Configuration Overview	502
	Configuring the LTE Mini-PIM as the Primary Interface	503
	Configuring the LTE Mini-PIM as a Backup Interface	505
	Configuring the LTE Interface as a Dial-on-Demand Interface	507

Part 8	Configuring Modem Interfaces	
Chapter 33	Configuring 3G Wireless Modems for WAN Connections	513
	3G Wireless Modem Overview	513
	3G Wireless Modem Configuration Overview	514
	Understanding the Dialer Interface	516
	Dialer Interface Configuration Rules	516
	Dialer Interface Authentication Support for GSM HSDPA 3G Wireless Modems	517
	Dialer Interface Functions	517
	Dialer Interface Operating Parameters	517
	Example: Configuring the Dialer Interface	518
	Understanding the 3G Wireless Modem Physical Interface	523
	Example: Configuring the 3G Wireless Modem Interface	524
	Understanding the GSM Profile	525
	Example: Configuring the GSM Profile	526
Chapter 34	Configuring CDMA EV-DO Modem Cards	529
	Understanding Account Activation for CDMA EV-DO Modem Cards	529
	Obtaining Electronic Serial Number (ESN)	529
	Account Activation Modes	530
	Activating the CDMA EV-DO Modem Card with IOTA Provisioning	531
	Activating the CDMA EV-DO Modem Card with OTASP Provisioning	532
	Activating the CDMA EV-DO Modem Card Manually	533
	Unlocking the GSM 3G Wireless Modem	535
Chapter 35	Configuring USB Modems for Dial Backup	537
	USB Modem Interface Overview	537
	USB Modem Interfaces	538
	Dialer Interface Rules	538
	How the Device Initializes USB Modems	539
	USB Modem Configuration Overview	540
	Example: Configuring a USB Modem Interface	542
	Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup	545
	Example: Configuring a Dialer Interface for USB Modem Dial-In	551
	Example: Configuring PAP on Dialer Interfaces	553
	Example: Configuring CHAP on Dialer Interfaces	554
Chapter 36	Configuring DOCSIS Mini-PIM Interfaces	557
	DOCSIS Mini-PIM Interface Overview	557
	Software Features Supported on DOCSIS Mini-PIMs	559
	Example: Configuring the DOCSIS Mini-PIM Interfaces	560
Chapter 37	Configuring Serial Interfaces	565
	Serial Interfaces Overview	565
	Serial Transmissions	566
	Signal Polarity	567
	Serial Clocking Modes	567
	Serial Interface Transmit Clock Inversion	568
	DTE Clock Rate Reduction	568

Serial Line Protocols	568
EIA-530	569
RS-232	569
RS-422/449	570
V.35	570
X.21	571
Example: Configuring a Serial Interface	571
Example: Deleting a Serial Interface	574
Understanding the 8-Port Synchronous Serial GPIM	575
Supported Features	575
Example: Configuring an 8-Port Synchronous Serial GPIM in Back-to-Back	
SRX650 Services Gateways	577

Part 9

Chapter 38

Configuration Statements and Operational Commands

Configuration Statements	595
accept-source-mac	598
access-point-name	599
apply-groups	599
arp-resp	600
authentication-method (Interfaces)	600
bandwidth (Interfaces)	601
bundle (Interfaces)	601
cbr rate	602
cellular-options	603
classifiers (CoS)	604
client-identifier (Interfaces)	605
code-points (CoS)	606
compression-device (Interfaces)	607
credit (Interfaces)	607
data-rate	608
disable (PoE)	609
dhcp (Interfaces)	610
duration (PoE)	611
family inet (Interfaces)	612
family inet6	615
flag (Interfaces)	618
flexible-vlan-tagging (Interfaces)	619
flow-control (Interfaces)	619
flow-monitoring (Services)	620
forwarding-classes (CoS)	621
fpc (Interfaces)	622
gratuitous-arp-reply	623
gsm-options	624
guard-band (PoE)	625
hub-assist	625
inline-jflow (Forwarding Options)	626
interface (PIC Bundle)	626
interface (PoE)	627

interfaces (CoS)	628
interval (Interfaces)	629
interval (PoE)	629
ipv4-template (Services)	630
ipv6-template (Services)	630
lACP (Interfaces)	631
latency (Interfaces)	631
lease-time	632
line-rate (Interfaces)	633
link-speed (Interfaces)	633
loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)	634
loss-priority (CoS Loss Priority)	635
loss-priority (CoS Rewrite Rules)	636
loss-priority-maps (CoS Interfaces)	637
loss-priority-maps (CoS)	637
management (PoE)	638
maximum-power (PoE)	639
media-type (Interfaces)	639
minimum-links (Interfaces)	640
mtu	641
native-vlan-id	642
next-hop-tunnel	644
no-dns-propagation	645
option-refresh-rate (Services)	645
pic-mode (Chassis TI Mode)	646
periodic (Interfaces)	647
ppp-over-ether	648
pppoe	649
pppoe-options	650
priority (PoE)	651
profile (Access)	652
profiles	654
promiscuous-mode (Interfaces)	655
quality (Interfaces)	655
r2cp	656
radio-router (Interfaces)	657
redundancy-group (Interfaces)	658
redundant-ether-options	659
redundant-parent (Interfaces Fast Ethernet)	660
redundant-parent (Interfaces Gigabit Ethernet)	661
request pppoe connect	662
request pppoe disconnect	663
resource (Interfaces)	664
retransmission-attempt	664
retransmission-interval (Interfaces)	665
roaming-mode	665
scheduler-map (CoS Virtual Channels)	666
select-profile	666
server-address	667

	shaping-rate (CoS Interfaces)	668
	simple-filter (Interfaces)	669
	sip-password	669
	sip-user-id	670
	source-address-filter (Interfaces)	671
	source-filtering (Interfaces)	672
	speed (Interfaces)	672
	speed (Gigabit Ethernet interface)	673
	speed (Chassis Cluster)	675
	telemetries (PoE)	676
	template-refresh-rate (Services)	677
	threshold (Interfaces)	677
	traceoptions (Interfaces)	678
	update-server	678
	vbr rate	679
	vdsl-profile	680
	vendor-id (Interfaces)	681
	web-authentication (Interfaces)	682
Chapter 39	Operational Commands	683
	clear oam ethernet connectivity-fault-management path-database	685
	clear dhcpv6 server binding (Local Server)	686
	clear ethernet-switching statistics mac-learning	687
	clear interfaces statistics swfabx	688
	clear ipv6 neighbors	689
	clear lacp statistics interfaces	690
	restart (Reset)	691
	request modem wireless create-profile	696
	request modem wireless fota	698
	request modem wireless sim-lock	699
	request modem wireless sim-unlock	700
	show chassis fpc (View)	702
	show chassis hardware (View)	710
	show ethernet-switching mac-learning-log	722
	show ethernet-switching table	727
	show igmp-snooping route (View)	750
	show interfaces	752
	show interfaces diagnostics optics	830
	show interfaces flow-statistics	835
	show interfaces queue	840
	show interfaces statistics (View)	843
	show interfaces terse zone	844
	show ipv6 neighbors	845
	show lacp interfaces (View)	847
	show lacp statistics interfaces (View)	851
	show modem wireless firmware	853
	show modem wireless network	856
	show modem wireless profiles	859
	show oam ethernet link-fault-management	861

show poe controller (View)	866
show pppoe interfaces	867
show pppoe statistics	871
show poe telemetries	873
show services accounting	875
show services accounting aggregation (View)	878
show services accounting aggregation template (View)	879
show services accounting flow-detail (View)	880

List of Figures

Part 1	Overview	
Chapter 2	Configuring Interface Logical Properties	21
	Figure 1: Subnets in a Network	24
Chapter 4	Configuring VLAN Tagging	55
	Figure 2: Typical LAN	56
	Figure 3: Typical VLAN	56
Part 2	Configuring DS1 Interfaces	
Chapter 7	Configuring DS3 Interfaces	81
	Figure 4: DS2 M-Frame Format	82
	Figure 5: DS3 M13 Frame Format	83
	Figure 6: DS3 C-Bit Parity Framing	84
Chapter 8	Configuring DS3 Interfaces	91
	Figure 7: DS2 M-Frame Format	92
	Figure 8: DS3 M13 Frame Format	93
	Figure 9: DS3 C-Bit Parity Framing	94
Part 3	Configuring DSL Interfaces	
Chapter 10	Configuring ADSL Interfaces	111
	Figure 10: MLPPP-over-ADSL Interface	125
Chapter 11	Configuring G.SHDSL Interfaces	145
	Figure 11: G.SHDSL Mini-PIM Operating in 2X4-Wire Mode	159
	Figure 12: G.SHDSL Mini-PIM Operating in 4X2-Wire Mode	159
	Figure 13: G.SHDSL Mini-PIM Operating in 1X8-Wire Mode	160
	Figure 14: G.SHDSL Mini-PIM Operating in EFM Mode	171
Chapter 12	Configuring VDSL2 Interfaces	181
	Figure 15: Typical VDSL2 End-to-End Connectivity and Topology Diagram	183
	Figure 16: Backward-Compatible ADSL Topology (ATM DSLAM)	183
	Figure 17: SRX Series Device with VDSL2 Mini-PIMs in an End-to-End Deployment Scenario	228
Part 4	Configuring Ethernet Interfaces	
Chapter 13	Performing Initial Configuration on Ethernet Interfaces	259
	Figure 18: Ethernet Frame Format	262
Chapter 15	Configuring Link Aggregation Control Protocol	295

	Figure 19: Topology for LAGs Connecting SRX Series Devices in Chassis Cluster to an EX Series Switch	304
Chapter 16	Configuring Gigabit Ethernet Physical Interface Modules	317
	Figure 20: Basic Back-to-Back Device Configuration	336
Chapter 18	Configuring Ethernet OAM Link Fault Management	355
	Figure 21: Ethernet LFM with SRX Series Devices	358
	Figure 22: Ethernet LFM with SRX Series Devices	362
Part 5	Configuring Interface Encapsulation	
Chapter 20	Interface Encapsulation Overview	381
	Figure 23: Frame Relay Network	382
Chapter 21	Configuring Point-to-Point Protocol over Ethernet	389
	Figure 24: PPPoE Session on the Ethernet Loop	400
	Figure 25: PPPoE Session on an ADSL Loop	401
	Figure 26: PPPoE Session on an ADSL Loop	404
Part 6	Configuring Link Services and Special Interfaces	
Chapter 24	Configuring Link Services Interfaces	431
	Figure 27: CRTP	435
	Figure 28: PPP and MLPPP Headers	448
Chapter 25	Configuring Link Fragmentation and Interleaving	455
	Figure 29: LFI on a Services Router	456
Chapter 27	Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles	471
	Figure 30: Configuring MLPPP and LFI on Serial Links	473
Part 7	Configuring LTE Interfaces	
Chapter 32	Configuring LTE Interfaces	499
	Figure 31: LTE Mini-PIM Used as a Primary Interface	503
	Figure 32: LTE Mini-PIM Used as a Backup Interface	505
	Figure 33: LTE Mini-PIM Used as a Dial-on-Demand Interface	508
Part 8	Configuring Modem Interfaces	
Chapter 33	Configuring 3G Wireless Modems for WAN Connections	513
	Figure 34: Wireless WAN Connections for Branch Offices	514
Chapter 36	Configuring DOCSIS Mini-PIM Interfaces	557
	Figure 35: Typical DOCSIS End-to-End Connectivity Diagram	558
Chapter 37	Configuring Serial Interfaces	565
	Figure 36: Serial Interface Clocking Modes	568
	Figure 37: Basic Back-to-Back Device Configuration	579

List of Tables

	About the Documentation	xxiii
	Table 1: Notice Icons	xxv
	Table 2: Text and Syntax Conventions	xxv
Part 1	Overview	
Chapter 1	Introduction to Interfaces	3
	Table 3: Network Interfaces	4
	Table 4: Configurable Services Interfaces	6
	Table 5: Non-Configurable Services Interfaces	7
	Table 6: Special Interfaces	8
	Table 7: Network Interface Names	9
	Table 8: IOC to NPC Connectivity Options	13
	Table 9: Port LED State for SFP+	19
	Table 10: Port LED State for QSFP	19
	Table 11: Port LED State for QSFP28	19
Chapter 2	Configuring Interface Logical Properties	21
	Table 12: Device Status Upon Configuration Change	31
Chapter 3	Understanding Interface Physical Properties	47
	Table 13: Interface Physical Properties	47
	Table 14: MTU Values for the SRX Series Services Gateways PIMs	51
Chapter 4	Configuring VLAN Tagging	55
	Table 15: VLAN ID Range by Interface Type Supported on the SRX Series Devices	57
	Table 16: Flexible VLANs	57
Chapter 5	Rate-Selectability	61
	Table 17: Rate Selectability of SRX4600 Services Gateway	63
	Table 18: Rate Selectability of SRX4600 Gateways	65
	Table 19: Maximum number of 10/40/100 Gigabit Ethernet ports Configurable at PIC and Port Level	65
	Table 20: Active Physical Ports on the SRX4600 Gateway for Configuring Rate Selectability at PIC level	69
	Table 21: Without number-of-ports But with Rate Selectability at PIC Level for SRX4600 Gateway	69
	Table 22: With number-of-ports Rate Selectability at PIC level for SRX4600 Gateway	70

Part 2	Configuring DS1 Interfaces	
Chapter 7	Configuring DS3 Interfaces	81
	Table 23: FEAC C-Bit Condition Indicators	85
Chapter 8	Configuring DS3 Interfaces	91
	Table 24: FEAC C-Bit Condition Indicators	95
Chapter 9	Configuring 1-Port Clear Channel DS3/E3 GPIM	101
	Table 25: 1-Port Clear Channel DS3/E3 GPIM Interface Options	103
Part 3	Configuring DSL Interfaces	
Chapter 10	Configuring ADSL Interfaces	111
	Table 26: Standard Bandwidths of DSL Operating Modes	111
Chapter 11	Configuring G.SHDSL Interfaces	145
	Table 27: Traffic Descriptors	147
	Table 28: Symmetrical WAN Speeds	148
	Table 29: Operating Wire Modes	160
	Table 30: Operating Wire Mode for EFM	171
Chapter 12	Configuring VDSL2 Interfaces	181
	Table 31: VDSL2 Annex A and Annex B Features	184
	Table 32: VDSL2 Operating Mode Backward Compatibility with ADSL	185
	Table 33: Supported Profiles on the VDSL2 Interfaces	186
Part 4	Configuring Ethernet Interfaces	
Chapter 13	Performing Initial Configuration on Ethernet Interfaces	259
	Table 34: Collision Backoff Algorithm Rounds	260
Chapter 18	Configuring Ethernet OAM Link Fault Management	355
	Table 35: Supported Interface Modes	356
Chapter 19	Configuring Power over Ethernet	367
	Table 36: PoE Specifications for the SRX210, SRX240, SRX320, SRX340, and SRX650 Devices	368
	Table 37: SRX Series Devices PoE Specifications	369
Part 6	Configuring Link Services and Special Interfaces	
Chapter 24	Configuring Link Services Interfaces	431
	Table 38: Services Available on a Link Services Interface	432
	Table 39: CoS Components Applied on Multilink Bundles and Constituent Links	443
	Table 40: PPP and MLPPP Encapsulation Overhead	448
	Table 41: Number of Packets Transmitted on a Queue	451
Chapter 26	Configuring Class-of-Service on Link Services Interfaces	459
	Table 42: Relative Priorities on Multilink Bundles and Constituent Links	463

Part 7	Configuring LTE Interfaces	
Chapter 32	Configuring LTE Interfaces	499
	Table 43: LTE Mini-PIM Models	499
Part 8	Configuring Modem Interfaces	
Chapter 35	Configuring USB Modems for Dial Backup	537
	Table 44: Default Modem Initialization Commands	539
	Table 45: Configuring Branch Office and Head Office Routers for USB Modem Backup Connectivity	541
	Table 46: Incoming Map Options	542
Chapter 36	Configuring DOCSIS Mini-PIM Interfaces	557
	Table 47: Software Features Supported on DOCSIS Mini-PIMs	559
Chapter 37	Configuring Serial Interfaces	565
	Table 48: Serial Transmission Signals	566
	Table 49: Supported Features	575
Part 9	Configuration Statements and Operational Commands	
Chapter 39	Operational Commands	683
	Table 50: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary	702
	Table 51: show chassis fpc Output Fields	704
	Table 52: show chassis hardware Output Fields	710
	Table 53: show ethernet-switching mac-learning-log Output Fields	722
	Table 54: show ethernet-switching mac-learning-log Output Fields	723
	Table 55: show ethernet-switching-mac-learning-log Output Fields	723
	Table 56: show ethernet-switching table Output Fields	730
	Table 57: show ethernet-switching table Output Fields	731
	Table 58: show ethernet-switching table Output fields	731
	Table 59: show ethernet-switching table Output Fields	732
	Table 60: show igmp-snooping route Output Fields	750
	Table 61: show interfaces (Gigabit Ethernet) Output Fields	756
	Table 62: Gigabit and 10 Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type	784
	Table 63: show interfaces Output Fields	785
	Table 64: show interfaces diagnostics optics Output Fields	830
	Table 65: show interfaces flow-statistics Output Fields	836
	Table 66: Flow Error Statistics (Packet Drop Statistics for the Flow Module)	836
	Table 67: show interfaces queue Output Fields	840
	Table 68: show ipv6 neighbors Output Fields	845
	Table 69: show lacp interfaces Output Fields	847
	Table 70: show lacp statistics interfaces Output Fields	851
	Table 71: show modem wireless firmware Output Fields	853
	Table 72: show modem wireless network Output Fields	856
	Table 73: show modem wireless profiles Output Fields	859
	Table 74: show oam ethernet link-fault-management Output Fields	861
	Table 75: show poe controller Output Fields	866

Table 76: show pppoe interfaces Output Fields	867
Table 77: show pppoe statistics Output Fields	871
Table 78: show poe telemetries interface Output Fields	873

About the Documentation

- Documentation and Release Notes on page xxiii
- Using the Examples in This Manual on page xxiii
- Documentation Conventions on page xxv
- Documentation Feedback on page xxvii
- Requesting Technical Support on page xxvii

Documentation and Release Notes

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

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

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

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

```
[edit]
```

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

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







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

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

Documentation Conventions

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

Table 1: Notice Icons

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

[Table 2 on page xxv](#) defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<code>user@host> show chassis alarms</code> <code>No alarms currently active</code>
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	<code>stub <default-metric metric>;</code>
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<code>broadcast multicast</code> <code>(string1 string2 string3)</code>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<code>rsvp { # Required for dynamic MPLS only</code>
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	<code>community name members [community-ids]</code>
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop address; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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



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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <https://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

PART 1

Overview

- [Introduction to Interfaces on page 3](#)
- [Configuring Interface Logical Properties on page 21](#)
- [Understanding Interface Physical Properties on page 47](#)
- [Configuring VLAN Tagging on page 55](#)
- [Rate-Selectability on page 61](#)

CHAPTER 1

Introduction to Interfaces

- [Understanding Interfaces on page 3](#)
- [Network Interfaces on page 4](#)
- [Services Interfaces on page 5](#)
- [Special Interfaces on page 8](#)
- [Interface Naming Conventions on page 8](#)
- [Understanding the Data Link Layer on page 10](#)
- [Configuring IOC to NPC Mapping on page 12](#)
- [Monitoring Interfaces on page 13](#)
- [Understanding GRE Keepalive Time on page 15](#)
- [Configuring GRE Keepalive Time on page 15](#)
- [SRX4600 LED Scheme Overview on page 19](#)

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.

- Related Documentation**
- [Interface Naming Conventions on page 8](#)
 - [Understanding Interface Logical Properties on page 21](#)
 - [Understanding Interface Physical Properties on page 47](#)
 - [Understanding the Data Link Layer on page 10](#)

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 3 on page 4](#) describes network interfaces that are available on SRX Series devices.

Table 3: Network Interfaces

Interface Name	Description
ae	Aggregated Ethernet interface. See “Understanding Aggregated Ethernet Interfaces” on page 279 .
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 523 and “LTE Mini-PIM Overview” on page 499 . 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” on page 537 and “LTE Mini-PIM Overview” on page 499 .
e1	E1 (also called DS1) WAN interface. See “Understanding T1 and E1 Interfaces” on page 73 .
e3	E3 (also called DS3) WAN interface. See “Understanding T3 and E3 Interfaces” on page 81 .
fe	Fast Ethernet interface. See “Understanding Ethernet Interfaces” on page 259 .
ge	Gigabit Ethernet interface. See “Understanding Ethernet Interfaces” on page 259 .
pt	VDSL2 interface. See “Example: Configuring VDSL2 Interfaces (Detail)” on page 227 .
reth	For chassis cluster configurations only, redundant Ethernet interface. See “Understanding Ethernet Interfaces” on page 259 .
se	Serial interface (either RS-232, RS-422/499, RS-530, V.35, or X.21). See “Serial Interfaces Overview” on page 565 .

Table 3: Network Interfaces (continued)

Interface Name	Description
t1	T1 (also called DS1) WAN interface. See “Understanding T1 and E1 Interfaces” on page 73.
t3	T3 (also called DS3) WAN interface. See “Understanding T3 and E3 Interfaces” on page 81.
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 325.



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

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.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Services Interfaces on page 5](#)
- [Special Interfaces on page 8](#)

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 4 on page 6](#) describes services interfaces that you can configure on SRX Series Services Gateways.

Table 4: Configurable Services Interfaces

Interface Name	Description
gr-0/0/0	<p>Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol inside another routing protocol.</p> <p>Packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then sent.</p> <p>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.</p> <p>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.</p>
ip-0/0/0	<p>Configurable IP-over-IP encapsulation (IP-IP tunnel) interface. IP tunneling allows the encapsulation of one IP packet inside another IP packet.</p> <p>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.</p> <p>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.</p> <p>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.</p>
lsq-0/0/0	<p>Configurable link services queuing interface. Link services include the multilink services MLPPP, MLFR, and Compressed Real-Time Transport Protocol (CRTP).</p> <p>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.</p> <p>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.</p>
lt-0/0/0	<p>Configurable logical tunnel interface that interconnects logical systems on SRX Series devices. See the Logical Systems Feature Guide for Security Devices.</p>
pp0	<p>Configurable PPPoE encapsulation interface. PPP packets being routed in an Ethernet network use PPPoE encapsulation.</p> <p>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.</p> <p>See “Understanding Point-to-Point Protocol over Ethernet” on page 389.</p>

Table 4: Configurable Services Interfaces (continued)

Interface Name	Description
ppd0	<p>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.</p> <p>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.</p> <p>Use the show pim interfaces command to check the status of ppd0 interface.</p>
ppe0	<p>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.</p> <p>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.</p>
st0	Secure tunnel interface used for IPSec VPNs. See the <i>IPsec VPN Feature Guide for Security Devices</i> .
umd0	<p>Configurable USB modem physical interface. This interface is detected when a USB modem is connected to the USB port on the device.</p> <p>See “USB Modem Configuration Overview” on page 540.</p>

Table 5 on page 7 describes non-configurable services interfaces for SRX Series Services Gateways.

Table 5: 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.

Table 5: Non-Configurable Services Interfaces (continued)

Interface Name	Description
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.

- Related Documentation**
- [Junos Services Interfaces Configuration](#)
 - [Understanding Interfaces on page 3](#)
 - [Network Interfaces on page 4](#)
 - [Special Interfaces on page 8](#)

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 6 on page 8](#) describes special interfaces for SRX Series Services Gateways.

Table 6: 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.

- Related Documentation**
- [Understanding Interfaces on page 3](#)
 - [Network Interfaces on page 4](#)
 - [Services Interfaces on page 5](#)

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 7 on page 9](#).

Table 7: 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.	<p>SRX5600 and SRX5800 devices: The slot number begins at 0 and increases as follows from left to right, bottom to top:</p> <ul style="list-style-type: none"> • SRX5600 device—Slots 0 to 5 • SRX5800 device—Slots 0 to 5, 7 to 11 <p>SRX3400 and SRX3600 devices: The Switch Fabric Board (SFB) is always 0. Slot numbers increase as follows from top to bottom, left to right:</p> <ul style="list-style-type: none"> • SRX3400 device—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.	<p>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.</p> <p>SRX3400, SRX3600, and SRX 4600 devices: This number is always 0. Only one IOC can be installed in a slot.</p>

Table 7: Network Interface Names (continued)

Name Part	Meaning	Possible Values
<i>port</i>	Number of the port on a PIM or IOC on which the physical interface is located.	<p>On SRX5600 and SRX5800 devices:</p> <ul style="list-style-type: none"> 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. <p>On SRX3400, SRX3600, and SRX 4600 devices:</p> <ul style="list-style-type: none"> 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. <p>Port numbers appear on the PIM or IOC faceplate.</p>
<i>channel</i>	Number of the channel (time slot) on a fractional or channelized T1 or E1 interface.	<ul style="list-style-type: none"> 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.
<i>unit</i>	Number of the logical interface created on a physical interface.	<p>A value from 0 through 16384.</p> <p>If no logical interface number is specified, unit 0 is the default, but must be explicitly configured.</p> <p>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.</p>



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

Related Documentation

- [Understanding Interfaces on page 3](#)

Understanding the Data Link Layer

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 on page 11](#)
- [Network Topology on page 11](#)

- [Error Notification on page 11](#)
- [Frame Sequencing on page 11](#)
- [Flow Control on page 11](#)
- [Data Link Sublayers on page 11](#)
- [MAC Addressing on page 12](#)

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**.

Related Documentation

- [Understanding Interfaces on page 3](#)

Configuring IOC to NPC Mapping

An Input/Output card (IOC) to Network Processing Card (NPC) mapping requires you to map one IOC to one NPC. However, you can map multiple IOCs to a single NPC. To balance the processing power in the NPC on the SRX3400, SRX3600, and SRX4600 Services Gateways, the chassis process (daemon) runs an algorithm that performs the mapping. It maps an IOC to an NPC that has the least amount of IOCs mapped to it. You can also use the command-line interface (CLI) to assign a specific IOC to a specific NPC.

When you configure the mapping, the chassis process will first use your configuration, then apply the least-number NPC algorithm for the rest of the IOCs.



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

To configure the IOC to NPC mapping:

```
[edit]
set chassis ioc-npc-connectivity {
  ioc slot-number npc (none | slot-number);
}
```

See [Table 8 on page 13](#) for a description of the **set chassis ioc-npc-connectivity** options.

Table 8: IOC to NPC Connectivity Options

Option	Description
<i>ioc slot-number</i>	Specify the IOC slot number. Range is 0 through 7 for SRX3400 devices and 0 through 12 for SRX3600 devices.
<i>npc slot-number</i>	Specify the NPC slot number. Range is 0 through 7 for SRX3400 devices and 0 through 12 for SRX3600 and SRX 4600 devices.
<i>none</i>	The chassis process maps the connection for the particular IOC.



NOTE: You must restart the chassis control after you commit the **set chassis ioc-npc-connectivity** command.

- Related Documentation**
- [Network Interfaces on page 4](#)
 - [Interface Naming Conventions on page 8](#)

Monitoring Interfaces

Purpose View general information about all physical and logical interfaces for a device.

Action Select **Monitor>Interfaces** in the J-Web user interface. The J-Web Interfaces page displays the following details about each device interface:

- Port—Indicates the interface name.
- Admin Status—Indicates whether the interface is enabled (Up) or disabled (Down).
- Link Status—Indicates whether the interface is linked (Up) or not linked (Down).
- Address—Indicates the IP address of the interface.
- Zone—Indicates whether the zone is an untrust zone or a trust zone.

- Services—Indicates services that are enabled on the device, such as HTTP and SSH.
- Protocols—Indicates protocols that are enabled on the device, such as BGP and IGMP.
- Input Rate graph—Displays interface bandwidth utilization. Input rates are shown in bytes per second.
- Output Rate graph—Displays interface bandwidth utilization. Output rates are shown in bytes per second.
- Error Counters chart—Displays input and output error counters in the form of a bar chart.
- Packet Counters chart—Displays the number of broadcast, unicast, and multicast packet counters in the form of a pie chart. (Packet counter charts are supported only for interfaces that support MAC statistics.)

To change the interface display, use the following options:

- Port for FPC—Controls the member for which information is displayed.
- Start/Stop button—Starts or stops monitoring the selected interfaces.
- Show Graph—Displays input and output packet counters and error counters in the form of charts.
- Pop-up button—Displays the interface graphs in a separate pop-up window.
- Details—Displays extensive statistics about the selected interface, including its general status, traffic information, IP address, I/O errors, class-of-service data, and statistics.
- Refresh Interval—Indicates the duration of time after which you want the data on the page to be refreshed.
- Clear Statistics—Clears the statistics for the selected interface.

Alternatively, you can enter the following **show** commands in the CLI to view interface status and traffic statistics:

- **show interfaces terse**



.....
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.
.....

- **show interfaces detail**
- **show interfaces extensive**
- **show interfaces *interface-name***

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, 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.

- Related Documentation**
- [Configuring GRE Keepalive Time](#)
 - [keepalive-time](#)
 - [hold-time](#)

Configuring GRE Keepalive Time

Keepalive times are only configurable for the ATM-over-ADSL interface, which is no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM starting in Junos OS Release 15.1X49-D10.

- [Configuring Keepalive Time and Hold time for a GRE Tunnel Interface on page 16](#)
- [Display GRE Keepalive Time Configuration on page 16](#)
- [Display Keepalive Time Information on a GRE Tunnel Interface on page 17](#)

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

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

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
```

Local	Remote	Interface	Admin	Link	Proto
200.1.3.1/24		gr-3/3/0.3	up	up	inet
		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
```

[illegible]

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: Sendbroadcast-pkt-to-re

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
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.

Related Documentation • [Understanding GRE Keepalive Time on page 15](#)

SRX4600 LED Scheme Overview

LEDs on the interface cards display the status of the ports. In SRX4600 gateways, there are four port LEDs per port. Each port provides an individual status LED with two states signaled by the color/LED state: OFF and SOLID GREEN.

The following port LED display modes are defined:

Table 9: Port LED State for SFP+

State	LED 1 – 4 (starting from left)
Default (power on, with or without Transceiver)	Off
Link Up	Solid Green for 10G
Link Down or Fault	Off
Admin Disable/Down	Off

Table 10: Port LED State for QSFP

State	LED 1 – 4 (starting from left)
Default (power on, with or without Transceiver)	Off
Link Up in 40G mode	LED-1 Solid Green
Link Up in 4x10G mode	Corresponding LED Solid Green
Link Down or Fault in 40G mode	Off
Link Down or Fault in 4x10G mode	Corresponding LED OFF
Admin Disable/Down in 40G mode	Off
Admin Disable/Down in 4x10G mode	Corresponding LED OFF

Table 11: Port LED State for QSFP28

State	LED 1 – 4 (starting from left)
Default (power on, with or without Transceiver)	Off
Link Up in 100G mode	LED-1 Solid Green
Link Down or Fault in 100G mode	Off

Table 11: Port LED State for QSFP28 (continued)

State	LED 1 – 4 (starting from left)
Admin Disable/Down in 100G mode	Off

- Related Documentation**
- [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on SRX4600 Gateway on page 67](#)

CHAPTER 2

Configuring Interface Logical Properties

- [Understanding Interface Logical Properties on page 21](#)
- [Understanding Protocol Families on page 22](#)
- [Understanding IPv4 Addressing on page 23](#)
- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 26](#)
- [Configuring the inet6 IPv6 Protocol Family on page 29](#)
- [Enabling Flow-Based Processing for IPv6 Traffic on page 30](#)
- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 32](#)
- [Understanding IPv6 Support VDSL2 Interfaces on page 41](#)
- [Example: Configuring the IPv6 Address on an ADSL Interface on page 41](#)
- [Understanding MAC Limiting on Layer 3 Routing Interfaces on page 44](#)

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

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Understanding Protocol Families on page 22](#)
- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 26](#)

- [Understanding Virtual LANs on page 55](#)

Understanding Protocol Families

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 on page 22](#)
- [Other Protocol Suites on page 22](#)

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.

Related Documentation

- [Understanding Interface Logical Properties on page 21](#)

Understanding IPv4 Addressing

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 on page 23](#)
- [IPv4 Dotted Decimal Notation on page 24](#)
- [IPv4 Subnetting on page 24](#)
- [IPv4 Variable-Length Subnet Masks on page 25](#)

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 network, each wire or ring requires its own network number and identifying subnet address.

Figure 1 on page 24 shows two subnets in a network.

Figure 1: Subnets in a Network




Figure 1 on page 24 shows three devices connected to one subnet and three more devices connected to a second subnet. Collectively, the six devices and two subnets make up

the larger network. In this example, the network is assigned the network prefix **192.14.0.0**, a class B address. Each device has an IP address that falls within this network prefix.

In addition to sharing a network prefix (the first two octets), the devices on each subnet share a third octet. 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 **192.14.126.0** and the beta subnet has the IP address **192.14.17.0**.

The subnet address **192.14.17.0** can be represented as follows in binary notation:

```
11000000 . 00001110 . 00010001 . xxxxxxxx
```

Because the first 24 bits in the 32-bit address identify the subnet, the last 8 bits are not significant. To indicate the subnet, the address is written as **192.14.17.0/24** (or just **192.14.17/24**). The **/24** is the subnet mask (sometimes shown as **255.255.255.0**).

IPv4 Variable-Length Subnet Masks

Traditionally, subnets were divided by address class. Subnets had either 8, 16, or 24 significant bits, corresponding to 2^8 , 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.

Related Documentation

- [Understanding Interface Logical Properties on page 21](#)
- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 26](#)

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

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.

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, 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.

- Related Documentation**
- *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;
    }
  }
}
```

Related Documentation

- [Understanding IPv6 Address Space, Addressing, Address Format, and Address Types on page 26](#)
- [Enabling Flow-Based Processing for IPv6 Traffic on page 30](#)

Enabling Flow-Based Processing for IPv6 Traffic

You have the following options for handling IPv6 traffic:

- Drop—Do not forward IPv6 packets. This is the default behavior.

- Packet-based forwarding—Do not create a session and process according to packet-based features only (includes firewall filters and class of service).
- Flow-based forwarding—Create a session and process according to packet-based features (including firewall filters and class of service) but also flow-based security features, such as screens and firewall security policy.

To enable flow-based processing for IPv6 traffic, modify the **mode** statement at the **[edit security forwarding-options family inet6]** hierarchy level:

```
security {
  forwarding-options {
    family {
      inet6 {
        mode flow-based;
      }
    }
  }
}
```

The following example shows the CLI commands you use to configure forwarding for IPv6 traffic:

```
[edit]
user@host# set security forwarding-options family inet6 mode ?
Possible completions:
  drop                Disable forwarding
  flow-based          Enable flow-based forwarding
  packet-based        Enable packet-based forwarding

[edit]
user@host# set security forwarding-options family inet6 mode flow-based
user@host# show security forwarding-options
family {
  inet6 {
    mode flow-based;
  }
}
```

If you change the forwarding option mode for IPv6, you might need to perform a reboot to initialize the configuration change. [Table 12 on page 31](#) summarizes device status upon configuration change.

Table 12: Device Status Upon Configuration Change

Configuration Change	Commit Warning	Reboot Required	Impact on Existing Traffic Before Reboot	Impact on New Traffic Before Reboot
Drop to flow-based	Yes	Yes	Dropped	Dropped
Drop to packet-based	No	No	Packet-based	Packet-based
Flow-based to packet-based	Yes	Yes	None	Flow sessions created
Flow-based to drop	Yes	Yes	None	Flow sessions created

Table 12: Device Status Upon Configuration Change (continued)

Configuration Change	Commit Warning	Reboot Required	Impact on Existing Traffic Before Reboot	Impact on New Traffic Before Reboot
Packet-based to flow-based	Yes	Yes	Packet-based	Packet-based
Packet-based to drop	No	No	Dropped	Dropped

- Related Documentation**
- [Understanding IPv6 Addressing](#)
 - [Configuring the inet6 IPv6 Protocol Family on page 29](#)

Configuring Flow Aggregation to Use Version 9 Flow Templates

Use of version 9 allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration.



NOTE: Version 9 requires that you install a services PIC, such as the Adaptive Services PIC or the Multiservices PIC, in the device. On MX Series routers, the Multiservices DPC fulfills this requirement.

The following sections contain additional information:

- [Configuring the Traffic to Be Sampled on page 32](#)
- [Configuring the Version 9 Template Properties on page 33](#)
- [Restrictions on page 34](#)
- [Fields Included in Each Template Type on page 35](#)
- [inet Sampling Behavior on page 37](#)
- [Verification on page 37](#)
- [Examples: Configuring Version 9 Flow Templates on page 38](#)

Configuring the Traffic to Be Sampled

To specify sampling of IPv4, IPv6, or peer AS billing traffic, include the appropriate configuration of the **family** statement at the [edit forwarding-options sampling input] hierarchy level:

```
[edit forwarding-options sampling input]
family (inet ) {
  max-packets-per-second number;
  rate number;
  run-length number;
}
```

You can include **family inet**.



NOTE: If you specify sampling for peer AS billing traffic, the **family** statement supports only IPv4 and IPv6 traffic (**inet**). Peer AS billing traffic is enabled only at the global instance hierarchy level and is not available for per Packet Forwarding Engine instances.

Configuring the Version 9 Template Properties

To define the version 9 templates, include the following statements at the **[edit services flow-monitoring version9]** hierarchy level:

```
[edit services flow-monitoring version9]
template name {
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  option-refresh-rate packets packets seconds seconds;
  template-refresh-rate packets packets seconds seconds;
  (ipv4-template (Services) | ipv6-template (Services) | mpls-ipv4-template |
   mpls-template | peer-as-billing-template) {
    label-position [ positions ];
  }
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the **template *name*** statement.
- You then specify each template for the appropriate type of traffic by including the **ipv4-template**, **ipv6-template**, **inet-ipv4-template**, **inet-template**, or **peer-as-billing-template**.
- If the template is used for **inet** traffic, you can also specify up to three label positions for the **inet** header label data by including the **label-position** statement; the default values are **[1 2 3]**.
- Within the template definition, you can optionally include values for the **flow-active-timeout** and **flow-inactive-timeout** statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds. Values you specify in template definitions override the global timeout values configured at the **[edit forwarding-options sampling output flow-server]** hierarchy level.



NOTE: In active flow monitoring, the flow-server records are exported after a time period that is a multiple of 60 seconds and greater than or equal to the configured active timeout value. For example, if the active timeout value is 90 seconds, the flow-server records are exported at 120-second intervals. If the active timeout value is 150 seconds, the flow-server records are exported at 180-second intervals, and so forth.

- You can also include settings for the **option-refresh-rate** and **template-refresh-rate** statements within a template definition. For both of these properties, you can include

a timer value (in seconds) or a packet count (in number of packets). For the **seconds** option, the default value is 60 and the range is from 10 through 600. For the **packets** option, the default value is 4800 and the range is from 1 through 480,000.

- To filter IPV6 traffic on a media interface, the following configuration is supported:

```
interfaces interface-name {
  unit 0 {
    family inet {
      sampling {
        input;
        output;
      }
    }
  }
}
```

Restrictions

The following restrictions apply to version 9 templates:

- You cannot apply the two different types of flow aggregation configuration (flow-server version 5/8 and flow aggregation version 9) at the same time.
- Flow export based on an **inet-ipv4** template assumes that the IPv4 header follows the inet header. In the case of Layer 2 VPNs, the packet on the provider router (P router) would look like this:

```
inet | Layer 2 Header | IPv4
```

In this case, **inet-ipv4** flows are not created on the PIC, because the IPv4 header does not directly follow the inet header. Packets are dropped on the PIC and are accounted as parser errors.

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC would be dropped until the synchronization process is complete.

On SRX300, SRX320, SRX340, SRX345, and SRX550HM devices, flow monitoring IPv6 version 9 has the following limitations:

- MPLS is not supported.
- User-defined version 9 templates are not supported.
- Routing Engine based flow monitoring version 9 is not supported.
- Flow monitoring and accounting are not supported in chassis cluster mode.
- Flow monitoring and accounting are not supported on an ae interface.
- J-Web for IPv6 sampled packets is not supported.

- SNMP queries for IPv6 sampled packets are not supported
- Flow monitoring can be configured in version 5, version 8, or version 9 export mode. Up to eight version 9 collectors are supported in export mode.
- Scope of accounting of IPv6 flow monitoring version 9 packets associated with pseudointerfaces (such as IRB, ML, LAG, VLAN, and GRE) is not supported.
- Creation of an SCTP session (parallel to TCP) between an exporter and a collector for gathering flow monitoring information is not supported.
- Maximum flow sessions that might be supported include:
 - A device with 1-GB RAM, such as an SRX320 device, might support up to 15,000 flow monitoring sessions at a time.
 - A device with 2-GB RAM, such as an SRX650 device, might support up to 59,900 flow monitoring sessions at a time.



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

- Routing Engine based flow monitoring V5 or V8 mode is mutually exclusive with inline flow monitoring V9.
- SRX5400, SRX5600, and SRX5800 do not support multiple collectors like SRX300, SRX320, SRX340, SRX345, and SRX550HM devices. Only one V9 collector per IPv4 or IPv6 is supported
- Flow aggregation for V9 export is not supported.
- Only UDP over IPv4 or IPv6 protocol can be used as the transport protocol.
- Only the standard IPv4 or IPv6 template is supported for exporting flow monitoring records.
- User-defined or special templates are not supported for exporting flow monitoring records.
- Chassis cluster is supported without flow monitoring session synchronization.

Fields Included in Each Template Type

The following fields are common to all template types:

- Input interface
- Output interface
- Number of bytes
- Number of packets
- Flow start time
- Flow end time

The IPv4 template includes the following specific fields:

- IPv4 Source Address
- IPv4 Destination Address
- L4 Source Port
- L4 Destination Port
- IPv4 TOS
- IPv4 Protocol
- ICMP type and code
- TCP Flags
- Application Id
- Direction
- IPv4 Next Hop Address

The IPv6 template includes the following specific fields:

- IPv6 Source Address and Mask
- IPv6 Destination Address and Mask
- L4 Source Port
- L4 Destination Port
- IPv6 TOS
- IPv6 Protocol
- TCP Flags
- Application Id
- Direction
- IP Protocol Version
- IPv6 Next Hop Address
- Egress Interface Information
- Source Autonomous System (AS) number
- Destination AS number

The inet template includes the following specific fields:

- inet Label #1
- inet Label #2
- inet Label #3

- inet EXP Information
- FEC IP Address

The inet-IPv4 template includes all the fields found in the IPv4 and inet templates.

The peer AS billing template includes the following specific fields:

- IPV4 Class of Service (TOS)
- Ingress Interface
- BGP IPV4 Next Hop Address
- BGP Peer Destination AS Number

inet Sampling Behavior

This section describes the behavior when inet sampling is used on egress interfaces in various scenarios (label pop or swap) on provider routers (P routers).

1. You configure inet sampling on an egress interface on the P router and configure an inet flow aggregation template. The route action is label *pop* because penultimate hop popping (PHP) is enabled.

Previously, IPv4 packets (only) would have been sent to the PIC for sampling even though you configured inet sampling. No flows should be created, with the result that the parser fails.

With the current capability of applying inet templates, inet flows are created.

2. As in the first case, you configure inet sampling on an egress interface on the P router and configure an inet flow aggregation template. The route action is label *swap* and the swapped label is 0 (explicit null).

The resulting behavior is that inet packets are sent to the PIC. The flow being sampled corresponds to the label before the swap.

3. You configure a Layer 3 VPN network, in which a customer edge router (CE-1) sends traffic to a provider edge router (PE-A), through the P router, to a similar provider edge router (PE-B) and customer edge router (CE-2) on the remote end.

The resulting behavior is that you cannot sample inet packets on the PE-A to P router link.

Verification

To verify the configuration properties, you can use the **show services accounting aggregation template template-name *name*** operational mode command.

All other **show services accounting** commands also support version 9 templates, except for **show services accounting flow-detail** and **show services accounting aggregation aggregation-type**.

Examples: Configuring Version 9 Flow Templates

The following is a sample version 9 template configuration:

```
services {
  flow-monitoring {
    version9 {
      template ip-template {
        flow-active-timeout 20;
        flow-inactive-timeout 120;
        ipv4-template;
      }
      template inet-template-1 {
        inet-template {
          label-position [1 3 4];
        }
      }
      template inet-ipv4-template-1 {
        inet-ipv4-template {
          label-position [1 5 7];
        }
      }
      template peer-as-billing-template-1 {
        peer-as-billing-template;
      }
    }
  }
}
```

The following is a sample firewall filter configuration for inet traffic:

```
firewall {
  family inet {
    filter inet_sample {
      term default {
        then {
          accept;
          sample;
        }
      }
    }
  }
}
```

The following sample configuration applies the inet sampling filter on a networking interface and configures the AS PIC to accept both IPv4 and inet traffic:

```
inline-jflows {
  at-0/1/1 {
    unit 0 {
      family inet {
        filter {
          input inet_sample;
        }
      }
    }
  }
}
```

```

    }
  }
  sp-7/0/0 {
    unit 0 {
      family inet;
      family inet;
    }
  }
}

```

The following example applies the inet version 9 template to the sampling output and sends it to the AS PIC:

```

forwarding-options {
  sampling {
    input {
      family inet {
        rate 1;
      }
    }
    output {
      flow-active-timeout 60;
      flow-inactive-timeout 30;
      flow-server 1.2.3.4 {
        port 2055;
        version9 {
          template inet-ipv4-template-1;
        }
      }
      inline-jflow sp-7/0/0 {
        source-address 1.1.1.1;
      }
    }
  }
}

```

The following is a sample firewall filter configuration for the peer AS billing traffic:

```

firewall {
  family inet {
    filter peer-as-filter {
      term 0 {
        from {
          destination-class dcu-1;
          inline-jflow ge-2/1/0;
          forwarding-class class-1;
        }
        then count count_team_0;
      }
    }
    term 1 {
      from {
        destination-class dcu-2;
        inline-jflow ge-2/1/0;
        forwarding-class class-1;
      }
    }
  }
}

```

```
        then count count_team_1;
    }
    term 2 {
        from {
            destination-class dcu-3;
            inline-jflow ge-2/1/0;
            forwarding-class class-1;
        }
        then count count_team_2;
    }
}
}
```

The following sample configuration applies the firewall filter as a filter attribute under the forwarding-options hierarchy for CoS-level data traffic usage information collection:

```
forwarding-options {
    family inet {
        filter output peer-as-filter;
    }
}
```

The following example applies the peer-as-billing version 9 template to enable sampling of traffic for billing purposes:

```
forwarding-options {
    sampling {
    }
    input {
        rate 1;
    }
    family inet {
        output {
            flow-server 10.209.15.58 {
                port 300;
                version9 {
                    template {
                        peer-as;
                    }
                }
            }
            inline-jflow sp-5/2/0 {
                source-address 2.3.4.5;
            }
        }
    }
}
family inet {
    filter {
        output peer-as-filter;
    }
}
```


- Related Documentation**
- [Understanding Interface Logical Properties on page 21](#)

Understanding IPv6 Support VDSL2 Interfaces

SRX300, SRX320, SRX340, SRX345, and SRX550HM devices support IPv6 on the following DSL encapsulations:

- ATM physical interface encapsulations
 - atm-pvc
 - ethernet-over-atm
- ATM logical interface encapsulations
 - atm-snap
 - atm-ppp-vc-mux
 - atm-nlpid
 - atm-cisco-nlpid
 - atm-ppp-llc
 - ether-over-atm-llc



NOTE: The encapsulation types atm-vc-mux and ppp-over-ether-over-atm-llc do not include IPv6 support.

To configure IPv6 addresses on DSL interfaces in ATM or PTM mode, include the family protocol type as **inet6** at the **[edit interfaces]** hierarchy level.

- Related Documentation**
- [Understanding Interface Logical Properties on page 21](#)

Example: Configuring the IPv6 Address on an ADSL Interface

This example shows how to configure the IPv6 address on an ADSL interface.

- [Requirements on page 41](#)
- [Overview on page 42](#)
- [Configuration on page 42](#)
- [Verification on page 43](#)

Requirements

Before you begin, configure network interfaces as necessary. See “[Understanding Ethernet Interfaces](#)” on page 259.

Overview

In this example, you specify the following configuration parameters:

- Encapsulation type: Ethernet over ATM on DSL logical interface
- ATM virtual path identifier (VPI): 2
- Encapsulation type: Ethernet over ATM on DSL logical interface
- Encapsulation type for the ATM-for-ADSL logical unit: Ethernet over ATM LLC
- ATM virtual channel (VCI): 2.118
- IPv6 address and prefix: 13:13::1/64

Configuration

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

```
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
```

Step-by-Step Procedure 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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-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 at-1/0/0
encapsulation ethernet-over-atm;
atm-options {
    vpi 2;
}
unit 0 {
    encapsulation ether-over-atm-llc;
    vci 2.118;
    family inet6 {
        address 13:13::1/64;
    }
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying ADSL Interface Properties

Purpose Verify that the ADSL interface properties are configured properly.

Action From operational mode, enter the **show ipv6 neighbors** command. The output shows a summary of interface information.

```
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
13:13::1          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
```

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

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Configuring the inet6 IPv6 Protocol Family on page 29](#)
- [show ipv6 neighbors on page 845](#)
- [clear ipv6 neighbors on page 689](#)

Understanding MAC Limiting on Layer 3 Routing Interfaces

- [Overview on page 44](#)
- [Limitations on page 46](#)

Overview

The MAC limiting feature provides a mechanism for limiting MAC addresses on devices that are connected to a Layer 3 routed Gigabit Ethernet (GE), Fast Ethernet (FE), or 10 Gigabit Ethernet (XE) interface. With MAC filters, you can allow traffic with specific source MAC. Software-based MAC limiting is supported. MAC limiting is applicable only on interfaces with plain Ethernet or VLAN tagged encapsulation.

Both the physical interface level **source-address-filter** and logical interface level **accept-source-mac** configurations are supported on SRX100, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, and SRX650 devices. (Platform support depends on the Junos OS release in your installation.) The following considerations apply when you configure the **source-address-filter** and **accept-source-mac** statements:

- If only the logical level **accept-source-mac** statement is configured, traffic from only those configured MAC addresses will be allowed on the logical interface.
- If only the physical interface level **source-address-filter** statement is configured, the physical interface's *allowed* MAC addresses are also considered the *allowed* addresses for all the logical interfaces belonging to the physical interface. Incoming packets from any other source MAC addresses are dropped.
- If the physical interface level **source-address-filter** is configured under **gigether-options** (or **fastether-options**) and **accept-source-mac** is configured for one or more of its logical interfaces or VLANs, the allowed list of addresses is a combination of MAC addresses specified in both the statements. For logical interfaces and VLANs where the **accept-source-mac** statement is not configured, the physical interface's *allowed* list of addresses is considered.

You can configure an interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the **source-address-filter** or **accept-source-mac** statements:

- **Logical level MAC filter configuration on an untagged interface**

```
ge-0/0/10 {  
  unit 0 {  
    accept-source-mac {  
      mac-address 00:22:33:44:55:66;  
      mac-address 00:26:88:e9:a3:01;  
    }  
    family inet {  
      address 60.60.60.1/24;  
    }  
  }  
}
```

- **Physical level MAC filter configuration on an untagged interface**

```

ge-0/0/10 {
  gigether-options {
    source-address-filter {
      00:55:55:55:55:66;
      00:26:88:e9:a3:01;
    }
  }
  unit 0 {
    family inet {
      address 60.60.60.1/24;
    }
  }
}

```

- Physical and logical level MAC filter configurations on a tagged interface

```

ge-0/0/10 {
  vlan-tagging;
  gigether-options {
    source-address-filter {
      00:26:88:e9:a3:01;
    }
  }
  unit 0 {
    vlan-id 40;
    accept-source-mac {
      mac-address 00:22:33:44:55:66;
    }
    family inet {
      address 40.40.40.1/24;
    }
  }
  unit 1 {
    vlan-id 60;
    accept-source-mac {
      mac-address 00:55:55:55:55:66;
    }
    family inet {
      address 60.60.60.1/24;
    }
  }
}

```



NOTE: On untagged Gigabit Ethernet interfaces, you must not configure the `source-address-filter` and the `accept-source-mac` statements simultaneously. If these statements are configured for the same interfaces at the same time, an error message appears. However, in the case of tagged VLANs, both these statements can be configured simultaneously, if no identical MAC addresses are specified.

Limitations

The following limitations apply to MAC limiting support on Layer 3 routed GE, FE, or XE interfaces:

- You can configure only 32 MAC addresses per device.
- Only software-based MAC filtering is supported. Software-based MAC filtering impacts performance. The performance impact is proportional to the number of MAC addresses configured.
- MAC- based policer or rate limiting is not supported.
- You cannot configure broadcast or multicast address in the source-address-filter statement.
- MAC filtering is not supported on Aggregated Ethernet (AE), Fabric Ethernet, Point-to-Point Protocol over Ethernet (PPPoE), Routed VLAN interface (RVI), or VLAN interfaces.

MAC filtering is not supported on chassis clusters.

Related Documentation

- [Understanding Interface Logical Properties on page 21](#)

CHAPTER 3

Understanding Interface Physical Properties

- [Understanding Interface Physical Properties on page 47](#)
- [Understanding Bit Error Rate Testing on page 48](#)
- [Understanding Interface Clocking on page 49](#)
- [Understanding Frame Check Sequences on page 50](#)
- [MTU Default and Maximum Values on page 51](#)
- [Understanding Jumbo Frames Support for Ethernet Interfaces on page 53](#)

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 13 on page 47](#) summarizes some key physical properties of device interfaces.

Table 13: 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 48 .
bert-period	Bit error rate test (BERT) time period over which bit errors are sampled. See “Understanding Bit Error Rate Testing” on page 48 .
chap	Challenge Handshake Authentication Protocol (CHAP). Specifying chap enables CHAP authentication on the interface. See “Understanding CHAP Authentication on a PPPoE Interface” on page 407 .

Table 13: Interface Physical Properties (continued)

Physical Property	Description
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 49 .
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 381 .
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. See “MTU Default and Maximum Values” on page 51 .
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 407 .
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.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Understanding Bit Error Rate Testing on page 48](#)
- [Understanding Interface Clocking on page 49](#)
- [Understanding Frame Check Sequences on page 50](#)
- [MTU Default and Maximum Values on page 51](#)

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.

Related Documentation • [Understanding Interface Physical Properties on page 47](#)

Understanding Interface Clocking

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 *clock jitter*. Long-term variations in the signal are known as *clock wander*.

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 on page 50](#)
- [Explicit Clocking Signal Transmission on page 50](#)

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 master and the other operates as a clock slave. The clock master internally generates a clock signal that is transmitted across the data link. The clock slave 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.

- Related Documentation**
- [Understanding Interface Physical Properties on page 47](#)

Understanding Frame Check Sequences

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 page 50](#)
- [Two-Dimensional Parity on page 51](#)

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.

Related Documentation

- [Understanding Interface Physical Properties on page 47](#)

MTU Default and Maximum Values

The MTU values are by default without any MTU configurations. If the MTU value is set, then the formula $\text{IFF MTU (IP MTU)} = \text{IFD MTU (Media MTU)} - \text{L2 Overhead}$ is applicable. See [Table 14 on page 51](#) 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 14 on page 51](#) lists MTU values for the SRX Series Services Gateways Physical Interface Modules (PIMs).

Table 14: 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

Table 14: MTU Values for the SRX Series Services Gateways PIMs (continued)

PIM	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP MTU (Bytes)
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
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
atm-vcmux	1514	1514	1514
atm-nlpid	1514	1514	1510

Table 14: 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	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
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

Related Documentation • [Understanding Interface Physical Properties on page 47](#)

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 51](#).

Related Documentation

- [MTU Default and Maximum Values on page 51](#)

CHAPTER 4


Configuring VLAN Tagging

- [Understanding Virtual LANs on page 55](#)
- [VLAN IDs and Ethernet Interface Types Supported on the SRX Series Devices on page 57](#)
- [Configuring VLAN Tagging on page 57](#)

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 56](#) 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 56](#) shows a typical VLAN topology.

Figure 3: Typical VLAN

Related Documentation

- [Understanding Interface Logical Properties on page 21](#)

- *MPLS Applications Feature Guide*

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

Table 15 on page 57 lists VLAN ID range by interface type supported on SRX Series devices:

Table 15: 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
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.

- Related Documentation**
- [Understanding Interface Physical Properties on page 47](#)

Configuring VLAN Tagging

You can configure SRX300, SRX320, SRX340, SRX345, 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 16 on page 57 for flexible VLANs.

Table 16: Flexible VLANs

Number of Tags	VLAN ID
0 (Untagged)	Native

Table 16: Flexible VLANs (continued)

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

This topic includes the following sections:

- [Configuring Single-Tag Framing on page 58](#)
- [Configuring Dual Tagging on page 58](#)
- [Configuring Mixed Tagging on page 58](#)
- [Configuring Mixed Tagging Support for Untagged Packets on page 59](#)

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, 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;

```

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;
  }
}
```

Related Documentation

- [Understanding Virtual LANs on page 55](#)

CHAPTER 5

Rate-Selectability

- [SRX4600 Gateway Rate-Selectability Overview on page 61](#)
- [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on SRX4600 Gateway on page 67](#)

SRX4600 Gateway Rate-Selectability Overview

The maximum amount of data that can be transmitted through a port at any given second either by a network device or by a component of the network device (such as a line card) is known as the port speed. Port speed is measured in kilobits per second (Kbps), gigabits per second (Gbps), and terabytes per second (Tbps). If a port can be configured to support both single and multiple speeds, the port is known as a rate-selectable port. Because the port is part of a network device (router or switch) or a network component (such as MPC, MIC) the component is known as a rate-selectable component. Rate selectability enables you to configure different port speeds at the port level or at the PIC level.

The SRX4600 has four rate-selectable ports (referred to as PIC 0 ports of FPC 1) that can be configured as 100-Gigabit Ethernet ports or 40-Gigabit Ethernet ports, or each port can be configured as four 10-Gigabit Ethernet ports (by using a breakout cable). The SRX4600 also has eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports of FPC 1). The four rate-selectable ports supports QSFP28/QSFP+ transceivers, whereas the eight 10-Gigabit Ethernet ports supports SFP+ transceivers. Knowing the exact details of the port speeds for the PICs helps you to choose the speeds to configure on the ports or on the PICs. You can view the port speeds of the PIC by executing *show chassis pic* command. For more information, see *SRX4600 Services Gateway Overview* and [“Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on SRX4600 Gateway” on page 67](#).



NOTE: By default, the two PICs of FPC 1, PIC 0 operate at 40-Gbps speed and PIC 1 operate at 10-Gbps speed.

The SRX4600 supports two types of rate selectability configuration options:

- **PIC Level Configuration:** To configure all ports to operate at the same speed, you configure rate selectability at the PIC level.

- **Port Level Configuration:** To configure different port speeds for each port, you configure rate selectability at the port level, in which case only the ports that are configured are enabled.

To configure all ports to operate at the same speed, configure rate selectability at the PIC level, in which case you cannot configure the speed of individual ports. To configure rate selectability at the PIC level, use the *pic-mode* statement and specify the port speed. To configure different port speeds for each port, configure rate selectability at the port level, in which case only the ports that are configured are enabled. To configure rate selectability at the port level, use the *speed* statement to specify the speed of individual ports.

The examples below show the sample CLI command output of the port speed capability of the 4-port PIC 0 with QSFP+ transceivers and the 8-port PIC 1 with SFP+ transceivers on the SRX4600.

```
user@router> show chassis pic fpc-slot 1 pic-slot 0
...
Port Speed Information:

  Port  Capable Port Speeds
  0     4x10GE, 40GE, 100GE
  1     4x10GE, 40GE, 100GE
  2     4x10GE, 40GE, 100GE
  3     4x10GE, 40GE, 100GE
...
user@router> show chassis pic fpc-slot 1 pic-slot 1
...
Port Speed Information:

  Port  Capable Port Speeds
  0     10GE
  1     10GE
  2     10GE
  3     10GE
  4     10GE
  5     10GE
  6     10GE
  7     10GE
...
```

[Table 17 on page 63](#) summarizes the rate selectability of the SRX4600 Services Gateway.

Table 17: Rate Selectability of SRX4600 Services Gateway

PIC	Port Number	Port Speed Supported
PIC 0 (FPC 1)	0–3	100-Gigabit Ethernet 40-Gigabit Ethernet 4x10-Gigabit Ethernet NOTE: <ul style="list-style-type: none"> Default port speed is 40 Gigabit Ethernet. You can configure one or all four 10-Gigabit Ethernet or 100-Gigabit Ethernet port operating at 10-Gigabit Ethernet mode to operate in 1-Gigabit Ethernet mode.
PIC 1 (FPC 1)	0–7	10 Gigabit Ethernet You can configure one or all 10-Gigabit Ethernet ports operating at 10-Gigabit Ethernet mode to operate in 1-Gigabit Ethernet mode.

The SRX4600 Services Gateway 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). Starting in Junos OS Release 18.1R1, SRX4600 supports 1-Gbps port speed on the default 10-Gbps ports on its 8-port PICs and on two dedicated chassis cluster control ports on the 4-port chassis cluster PICs. 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.

1-Gbps speed is only supported in non-autonegotiation mode.

To configure 1-Gbps port speed:

- On the 8-port 10-Gigabit Ethernet PIC, use the **set interfaces *intf-name* gigether-options speed 1g** command and commit the configuration. Refer [speed \(Gigabit Ethernet interface\)](#) for more details.
- On chassis cluster control interfaces (on a 4-port 10-Gigabit Ethernet PIC), use the **set chassis cluster control-port speed 1g** command. Refer [speed \(Chassis Cluster\)](#) for more details.

You must reboot the device for the changed configuration to take effect.

**NOTE:**

- The interface name prefix must be xe.
- Only the eight ports on PIC 1 of FPC 1 support 1-Gigabit Ethernet mode. The ports on PIC 0 of FPC 1 only support 100-Gigabit Ethernet, 40-Gigabit Ethernet, 4x10-Gigabit Ethernet mode and not 1-Gigabit Ethernet mode.
- The rate selectability at PIC level and port level does not support 1-Gbps speed. But you can configure the port configured at 10-Gbps speed to operate at 1-Gbps speed at Gigabit Ethernet interface level using [speed \(Gigabit Ethernet interface\)](#) and [speed \(Chassis Cluster\)](#) configuration statements respectively.
- The 1-Gbps operation mode is only supported in non-autonegotiation mode.

To view the speed configured for the interface, execute the **show interfaces extensive** command. The **Speed Configuration** output parameter in the command output indicates the current operation speed of the interface. If the interface is configured with 1-Gbps speed, then **Speed Configuration** displays **1G**; if the interface is configured with 10-Gbps speed, **Speed Configuration** displays **AUTO**.

For example:

```
user@router > show interfaces xe-1/1/0 extensive
Physical interface: xe-1/1/0, Enabled, Physical link is Down
Interface index: 151, SNMP ifIndex: 613, Generation: 154
Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 10Gbps,
BPDU Error: None, Loop Detect PDU Error: None, MAC-REWRITE Error: None,
Loopback: None, Source filtering: Disabled, Flow control: Enabled,
Speed Configuration: 1G
...
```

In this example, the **Speed Configuration** output parameter displays **1G**, which means the operation speed of xe-1/1/0 interface is 1-Gbps speed.

User-Configurable Rate Selectability of SRX4600 Services Gateway

You can also configure rate selectability on SRX4600 gateways.



NOTE: Only the Interface that is already operating at 10-Gigabit Ethernet mode can be configured to operate at 1-Gigabit Ethernet mode.

Table 18 on page 65 summarizes the user-configurable rate selectability of SRX4600 gateways.

Table 18: Rate Selectability of SRX4600 Gateways

Port Speed Configuration on PIC 0 (Gbps)	Port Speed Configuration on PIC 1 (Gbps)
100	0 Configure the number of active ports to 0
10	10
40	0 Configure the number of active ports to 0



NOTE: The SRX4600 does not support heterogeneous mode. That is, in PIC mode if 40-Gbps or 100-Gbps speed is configured on PIC 0, then the *number-of-ports* on PIC 1 must be configured to 0 only.

Maximum number of 10/40/100GE ports Configurable at PIC and Port Mode

Following table summarizes the maximum number of 10/40/100 Gigabit Ethernet ports per PIC configurable at PIC and port levels:

Table 19: Maximum number of 10/40/100 Gigabit Ethernet ports Configurable at PIC and Port Level

Maximum Ports	Maximum Ports configurable at PIC Mode (on both PICO and PIC1)	Maximum Ports Configurable at Port Mode (on both PICO and PIC1)
10 Gigabit Ethernet Ports	24 Which means 16 ports from PIC 0 and 8 Ports from PIC 1.	20 Which means 12 ports from PIC 0 and eight ports from PIC 1.
40 Gigabit Ethernet Ports	4 Only four ports from PIC 0 as PIC 1 supports only 10 Gbps Speed.	4
100 Gigabit Ethernet Ports	4 Only four ports from PIC 0 as PIC 1 supports only 10 Gbps speed.	4

Port Configuration - PIC Level

On PIC 0, if each of the four ports is configured to operate at 100-Gbps speed, then you must configure all the 8 ports at PIC 1 to 0 (using *number-of-ports* statement). On PIC 0, if ports 0, 1, and 2 are set to 100-Gbps, and port 3 is set to 10-Gbps or 40-Gbps, then you should configure all the 8 ports at PIC 1 to 0 (using **number-of-ports** statement), and so on



NOTE: Only the Interface that is already operating at 10GE mode can be configured to operate at 1GE mode.

Configuring 40-Gigabit Ethernet ports to 4X10-Gigabit Ethernet using Breakout Cables

On SRX4600 Services Gateway, Slot 1 PIC 0 comes up with the default setting of 4X40-Gigabit Ethernet ports. Starting in Junos OS Release 18.1R1, you can configure the 40 Gigabit Ethernet port to 4X10-Gigabit Ethernet mode by plugging in QSFP-4X10-Gigabit Ethernet optics connecting with 4x10-Gigabit Ethernet breakout cables. You use QSFP+ transceivers to connect the 40-Gbps (default speed) port to the breakout cable, which connects to four SFP+ transceivers at the other end thus converting that port into four 10-Gbps interfaces).

In 4X10-Gigabit Ethernet mode, the interface naming convention for 10-Gigabit Ethernet interface is as follows:

Interface Type	Interface Convention (MPC Slot: x, PIC Slot: y, Port: z)
4x10-Gigabit Ethernet	xe-x/y/z:0
	xe-x/y/z:1
	xe-x/y/z:2
	xe-x/y/z:3
	For example, when 40-Gigabit Ethernet port et-1/0/0 is converted to 4x10-Gigabit Ethernet ports, the naming convention of 10-Gigabit Ethernet port is:
	xe-1/0/0:0
	xe-1/0/0:1
	xe-1/0/0:2
	xe-1/0/0:3

**NOTE:**

- The 40-Gigabit Ethernet ports that are converted to four 10-Gigabit Ethernet ports cannot be configured to 1-Gbps speed.
- The SRX4600 Gateway supports 40-Gigabit Ethernet breakouts only in PIC mode.

For example:

In PIC mode, to configure all 40-Gigabit Ethernet ports of PIC 0, execute the following configuration statement:

```
set chassis fpc 1 pic 0 pic-mode 10G
```

In PIC mode, to set only the first two 40-Gigabit Ethernet ports, execute the following configuration statement:

```
set chassis fpc 1 pic 0 pic-mode 10G number-of-ports 2
```

This configuration sets only the first two ports of 40-Gigabit Ethernet port and leaves the last two ports disabled, that means the last two ports cannot be used as 40-Gigabit Ethernet ports after this configuration.

After you commit the configuration, for the new configuration to take effect, you must reboot the device or chassis cluster.

Release History Table

Release	Description
18.1R1	Starting in Junos OS Release 18.1R1, SRX4600 supports 1-Gbps port speed on the default 10-Gbps ports on its 8-port PICs and on two dedicated chassis cluster control ports on the 4-port chassis cluster PICs.

Related Documentation

- [Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on SRX4600 Gateway on page 67](#)
- *speed*
- *show chassis pic*
- *number-of-ports*
- *pic-mode*

Supported Active Physical Rate-Selectable Ports to Prevent Oversubscription on SRX4600 Gateway

The maximum capacity of an SRX4600 gateway is 240 Gbps, which cannot be oversubscribed. In SRX4600, the network ports are available in two groups (referred to as PICs), with restrictions around the number and type of ports that can be configured without oversubscription.

Starting in Junos OS Release 17.4R1, SRX4600 gateways support rate selectability to prevent oversubscription of the Packet Forwarding Engine bandwidth. The SRX4600 Packet Forwarding Engine has four 100-Gigabit Ethernet QSFP28 ports (referred to as PIC 0 ports) and eight 10-Gigabit Ethernet ports (referred to as PIC 1 ports). Each of the PIC 0 ports can be used as either a 100-Gigabit Ethernet QSFP28 port or a 40-Gigabit Ethernet QSFP port, or they can be configured as four 10-Gigabit Ethernet ports (using a breakout cable).

If you configure rate selectability at the PIC level, all the ports supporting that port speed are enabled by default. This can lead to oversubscription in certain cases. To prevent the oversubscription, you can configure the number of active ports that operate at the configured speed by using the **`number-of-ports number-of-active-physical-ports`** configuration statement. Additionally, interfaces are created only for the active ports.

**NOTE:**

- You cannot configure the number of active ports when you configure rate selectability at the port level.
-

Invalid Port Configuration

You must try to avoid configuring ports that can lead to oversubscription.

Following is an example of an invalid configuration:

```
4x100GE + 8x10GE
```

If you try to commit an invalid configuration, the configuration gets committed, but the port is not activated. This is because Junos OS allows you to configure a port before a line card is inserted. You will get an error message in the output of the `show chassis alarms` command and also in the log messages.

Configuring Active Ports on SRX4600 Gateway with Rate Selectability

Table 20 on page 69 summarizes the active ports with *number-of-ports* configured but without any rate selectability configuration for an SRX4600 gateway. Because there is no rate selectability configured, the default speed is used in these cases.

Table 20: Active Physical Ports on the SRX4600 Gateway for Configuring Rate Selectability at PIC level

PIC	Number of Ports (<i>number-of-ports</i> Statement)	Active Ports		
		PIC Level 10-Gigabit Ethernet Profile	PIC Level 40-Gigabit Ethernet Profile	PIC Level 100-Gigabit Ethernet Profile
PIC 0	0	-	-	-
	1	0	-	-
	2	0, 1	-	-
	3	0, 1, 2	-	-
	4	0, 1, 2, 3	-	-
PIC 1	0	-	-	-
	1	0	-	-
	2	0, 1	-	-
	3	0, 1, 2	-	-
	4	0, 1, 2, 3	-	-
	5	0, 1, 2, 3, 4	-	-
	6	0, 1, 2, 3, 4, 5	-	-
	7	0, 1, 2, 3, 4, 5, 6	-	-
	8	0, 1, 2, 3, 4, 5, 6, 7	-	-

[Table 21 on page 69](#) summarizes the active ports without *number-of-ports* configured but with rate selectability at PIC-level configuration for an SRX4600 gateway.

Table 21: Without *number-of-ports* But with Rate Selectability at PIC Level for SRX4600 Gateway

PIC	Active Ports		
	PIC-Level 10-Gigabit Ethernet	PIC-Level 40-Gigabit Ethernet	PIC-Level 100-Gigabit Ethernet
PIC 0	0-3	0-3	0-3
PIC 1	0-7	-	-

[Table 22 on page 70](#) summarizes the active ports with *number-of-ports* configured and rate selectability at PIC-level configuration for an SRX4600 gateway.

Table 22: With number-of-ports Rate Selectability at PIC level for SRX4600 Gateway

PIC	Number of Ports (number-of-ports Statement)	Active Ports		
		PIC-Level 10-Gigabit Ethernet	PIC-Level 40-Gigabit Ethernet	PIC-Level 100-Gigabit Ethernet
PIC 0	0	-	-	-
	1	0	0	0
	2	0, 1	0, 1	0, 1
	3	0, 1, 2	0, 1, 2	0, 1, 2
	4	0, 1, 2, 3	0, 1, 2, 3	0, 1, 2, 3
PIC 1	0	-	-	-
	1	0	-	-
	2	0, 1	-	-
	3	0, 1, 2	-	-
	4	0, 1, 2, 3	-	-
	5	0, 1, 2, 3, 4	-	-
	6	0, 1, 2, 3, 4, 5	-	-
	7	0, 1, 2, 3, 4, 5, 6	-	-
	8	0, 1, 2, 3, 4, 5, 6, 7	-	-

Related Documentation • [SRX4600 LED Scheme Overview on page 19](#)

PART 2

Configuring DS1 Interfaces

- [Configuring DS1 Interfaces on page 73](#)
- [Configuring DS3 Interfaces on page 81](#)
- [Configuring DS3 Interfaces on page 91](#)
- [Configuring 1-Port Clear Channel DS3/E3 GPIM on page 101](#)

CHAPTER 6

Configuring DS1 Interfaces

- [Understanding T1 and E1 Interfaces on page 73](#)
- [Example: Configuring a T1 Interface on page 76](#)
- [Example: Deleting a T1 Interface on page 79](#)

Understanding T1 and E1 Interfaces

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 on page 73](#)
- [E1 Overview on page 74](#)
- [T1 and E1 Signals on page 74](#)
- [Encoding on page 74](#)
- [T1 and E1 Framing on page 75](#)
- [T1 and E1 Loopback Signals on page 75](#)

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 193 bits per T1 frame. Frames are transmitted 8,000 times per second, at a data transmission rate of 1.544 Mbps ($8,000 \times 193 = 1.544$ 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 use 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.

**Related
Documentation**

- [Example: Configuring a T1 Interface on page 76](#)

Example: Configuring a T1 Interface

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

- [Requirements on page 76](#)
- [Overview on page 76](#)
- [Configuration on page 76](#)
- [Verification on page 77](#)

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

**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

Confirm that the configuration is working properly.

- [Verifying the Link State of All Interfaces on page 77](#)
- [Verifying Interface Properties on page 78](#)

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 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.

Related Documentation

- [Understanding T1 and E1 Interfaces on page 73](#)
- [Example: Deleting a T1 Interface on page 79](#)

Example: Deleting a T1 Interface

This example shows how to delete a T1 interface.

- [Requirements on page 79](#)
- [Overview on page 79](#)
- [Configuration on page 79](#)
- [Verification on page 79](#)

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

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.

Related Documentation

- [Understanding T1 and E1 Interfaces on page 73](#)
- [Example: Configuring a T1 Interface on page 76](#)

CHAPTER 7

Configuring DS3 Interfaces

- [Understanding T3 and E3 Interfaces on page 81](#)
- [Example: Configuring a T3 Interface on page 86](#)
- [Example: Deleting a T3 Interface on page 88](#)

Understanding T3 and E3 Interfaces

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 on page 81](#)
- [DS2 Bit Stuffing on page 82](#)
- [DS3 Framing on page 82](#)

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 82](#) 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 83](#) and [Figure 6 on page 84](#).

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 83](#).

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.
- Multiframe 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 pattern for each M-frame is 010.
- Bit stuffing control bits (C-bits)—Serve as bit stuffing indicators for each DS2 input. For example, C₁₁, C₁₂, and C₁₃ 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 84](#).

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 **0xxxxxx 1111111**, in which x can be either 1 or 0. Bits are transmitted from right to left.

[Table 23 on page 85](#) lists some C-bit code words and the alarm or status condition indicated.

Table 23: FEAC C-Bit Condition Indicators

Alarm or Status Condition	C-Bit Code Word
DS3 equipment failure requires immediate attention.	00110010 1111111
DS3 equipment failure occurred—such as suspended, not activated, or unavailable service—that is non-service-affecting.	00011110 1111111
DS3 loss of signal.	00011100 1111111
DS3 out of frame.	00000000 1111111
DS3 alarm indication signal (AIS) received.	00101100 1111111
DS3 idle received.	00110100 1111111
Common equipment failure occurred that is non-service-affecting.	00011101 1111111
Multiple DS1 loss of signal.	00101010 1111111
DS1 equipment failure occurred that requires immediate attention.	00001010 1111111
DS1 equipment failure occurred that is non-service-affecting.	00000110 1111111
Single DS1 loss of signal.	00111100 1111111

- **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.

Related Documentation

- [Example: Configuring a T3 Interface on page 86](#)
- [Example: Deleting a T3 Interface on page 88](#)

Example: Configuring a T3 Interface

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

- [Requirements on page 86](#)
- [Overview on page 86](#)
- [Configuration on page 86](#)
- [Verification on page 87](#)

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

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

Confirm that the configuration is working properly.

- [Verifying the Link State of All Interfaces on page 87](#)
- [Verifying Interface Properties on page 88](#)

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.

Related Documentation

- [Understanding T3 and E3 Interfaces on page 81](#)
- [Example: Deleting a T3 Interface on page 88](#)

Example: Deleting a T3 Interface

This example shows how to delete a T3 interface.

- [Requirements on page 88](#)
- [Overview on page 89](#)
- [Configuration on page 89](#)
- [Verification on page 89](#)

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

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.

Related Documentation

- [Understanding T3 and E3 Interfaces on page 81](#)
- [Example: Configuring a T3 Interface on page 86](#)

CHAPTER 8

Configuring DS3 Interfaces

- [Understanding T3 and E3 Interfaces on page 91](#)
- [Example: Configuring a T3 Interface on page 96](#)
- [Example: Deleting a T3 Interface on page 98](#)

Understanding T3 and E3 Interfaces

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.

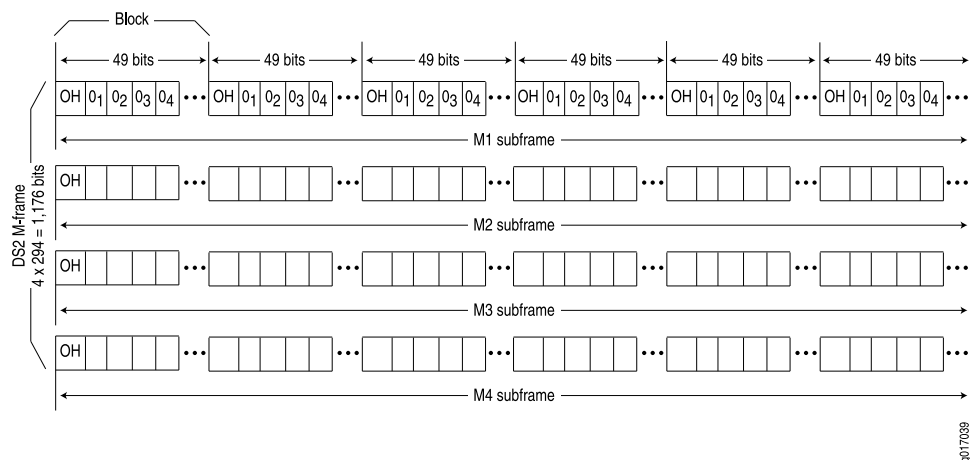
- [Multiplexing DS1 Signals on page 91](#)
- [DS2 Bit Stuffing on page 92](#)
- [DS3 Framing on page 92](#)

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 82](#) shows the DS2 M-frame format.

Figure 7: 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 O_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.

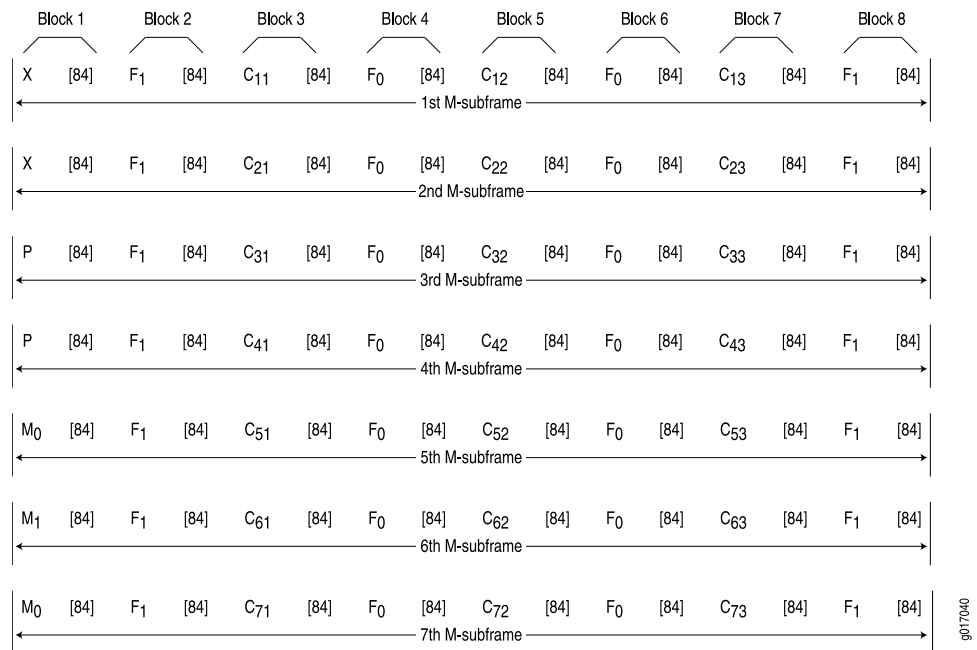
DS3 Framing

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 83](#) and [Figure 6 on page 84](#).

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 83](#).

Figure 8: 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 pattern for each M-frame is 010.
- **Bit stuffing control bits (C-bits)**—Serve as bit stuffing indicators for each DS2 input. For example, C₁₁, C₁₂, and C₁₃ 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

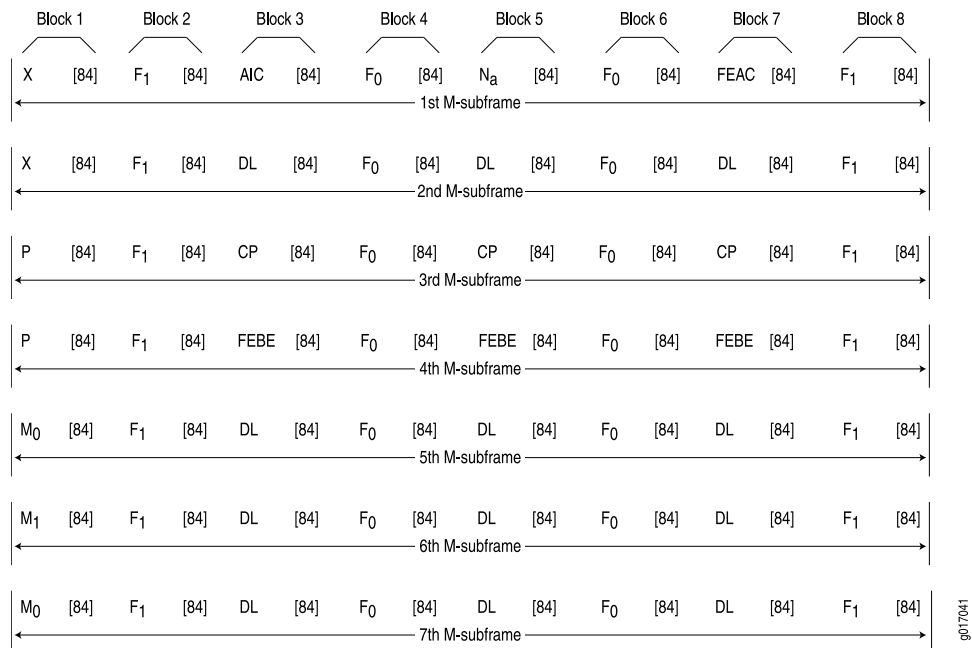
side, the number of 1s for a given frame does not match the P-bits in the following frame, it indicates one or more bit errors in the transmission.

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 84](#).

Figure 9: 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 **0xxxxxx 1111111**, in which x can be either 1 or 0. Bits are transmitted from right to left.

[Table 23 on page 85](#) lists some C-bit code words and the alarm or status condition indicated.

Table 24: FEAC C-Bit Condition Indicators

Alarm or Status Condition	C-Bit Code Word
DS3 equipment failure requires immediate attention.	00110010 1111111
DS3 equipment failure occurred—such as suspended, not activated, or unavailable service—that is non-service-affecting.	00011110 1111111
DS3 loss of signal.	00011100 1111111
DS3 out of frame.	00000000 1111111
DS3 alarm indication signal (AIS) received.	00101100 1111111
DS3 idle received.	00110100 1111111
Common equipment failure occurred that is non-service-affecting.	00011101 1111111
Multiple DS1 loss of signal.	00101010 1111111
DS1 equipment failure occurred that requires immediate attention.	00001010 1111111
DS1 equipment failure occurred that is non-service-affecting.	00000110 1111111
Single DS1 loss of signal.	00111100 1111111

- **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.

Related Documentation

- [Example: Configuring a T3 Interface on page 86](#)
- [Example: Deleting a T3 Interface on page 88](#)

Example: Configuring a T3 Interface

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

- [Requirements on page 96](#)
- [Overview on page 96](#)
- [Configuration on page 96](#)
- [Verification on page 97](#)

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

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

Confirm that the configuration is working properly.

- [Verifying the Link State of All Interfaces on page 97](#)
- [Verifying Interface Properties on page 98](#)

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.

Related Documentation

- [Understanding T3 and E3 Interfaces on page 81](#)
- [Example: Deleting a T3 Interface on page 88](#)

Example: Deleting a T3 Interface

This example shows how to delete a T3 interface.

- [Requirements on page 98](#)
- [Overview on page 99](#)
- [Configuration on page 99](#)
- [Verification on page 99](#)

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

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.

Related Documentation

- [Understanding T3 and E3 Interfaces on page 81](#)
- [Example: Configuring a T3 Interface on page 86](#)

CHAPTER 9

Configuring 1-Port Clear Channel DS3/E3 GPIM

- [Understanding the 1-Port Clear Channel DS3/E3 GPIM on page 101](#)
- [Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for M23 Mapping Mode on page 104](#)
- [Example: Configuring the 1-Port Clear-Channel DS3/E3 GPIM for DS3 Port Mode on page 106](#)
- [Example: Configuring the 1-Port Clear Channel DS3/E3 GPIM for E3 Port Mode on page 107](#)

Understanding the 1-Port Clear Channel DS3/E3 GPIM

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.



NOTE: Starting in Junos OS Release 15.1X49-D10, the 1-Port Clear Channel DS3/E3 interface is no longer supported on SRX650 devices.

This topic includes the following sections:

- [Supported Features on page 101](#)
- [Interface Naming on page 102](#)
- [Physical Interface Settings on page 102](#)
- [Logical Interface Settings on page 102](#)

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 25 on page 103 identifies network interface specifications for DS3 or E3 modes.

Table 25: 1-Port Clear Channel DS3/E3 GPIM Interface Options

Description	DS3 Mode	E3 Mode
Network Interface Specifications		
Line encoding	B3ZS	HDB3
Framing	<ul style="list-style-type: none"> C-bit parity (default) M23 	G.751 (default)
Subrate and scrambling	Vendor algorithms supported: <ul style="list-style-type: none"> Adtran Digital Link Kentrox Larscom Verilink 	Vendor algorithms supported: <ul style="list-style-type: none"> Digital Link Kentrox
Network alarms	Supported in accordance with the ANSI specification: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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)

Table 25: 1-Port Clear Channel DS3/E3 GPIM Interface Options (continued)

Description	DS3 Mode	E3 Mode
Shared flag	Supported	Supported
Idle flag/fill (0x7e or all ones)	Supported	Supported
Counters	Runts, giants	Runts, giants

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10, the 1-Port Clear Channel DS3/E3 interface is no longer supported on SRX650 devices.

Related Documentation

- [Interface Naming Conventions on page 8](#)

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

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



NOTE: Starting in Junos OS Release 15.1X49-D10, the 1-Port Clear Channel DS3/E3 interface is no longer supported on SRX650 devices.

- [Requirements on page 104](#)
- [Overview on page 104](#)
- [Configuration on page 105](#)

Requirements

Before you begin:

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

Overview

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

Configuration

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
```

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10, the 1-Port Clear Channel DS3/E3 interface is no longer supported on SRX650 devices.

Related Documentation

- [Understanding the 1-Port Clear Channel DS3/E3 GPIM on page 101](#)

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

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



NOTE: Starting in Junos OS Release 15.1X49-D10, the 1-Port Clear Channel DS3/E3 interface is no longer supported on SRX650 devices.

- [Requirements on page 106](#)
- [Overview on page 106](#)
- [Configuration on page 106](#)

Requirements

Before you begin:

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

Overview

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

Configuration

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
```

- 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
```

- Set the MTU value to 9018.

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

- Set the framing mode.

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

- Enable the unframed DS3 mode.

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

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

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

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

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

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10, the 1-Port Clear Channel DS3/E3 interface is no longer supported on SRX650 devices.

Related Documentation

- [Understanding the 1-Port Clear Channel DS3/E3 GPIM on page 101](#)

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

This example modifies the default configuration for an E3 environment.

- [Requirements on page 107](#)
- [Overview on page 108](#)
- [Configuration on page 108](#)

Requirements

Before you begin:

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

Overview

This example configures the basic E3 interface.

Configuration

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

Related Documentation

- [Understanding the 1-Port Clear Channel DS3/E3 GPIM on page 101](#)

PART 3

Configuring DSL Interfaces

- [Configuring ADSL Interfaces on page 111](#)
- [Configuring G.SHDSL Interfaces on page 145](#)
- [Configuring VDSL2 Interfaces on page 181](#)

Configuring ADSL Interfaces

- [ADSL Interface Overview on page 111](#)
- [ADSL and SHDSL Interfaces Configuration Overview on page 114](#)
- [Example: Configuring ATM-over-SHDSL Network Interfaces on page 118](#)
- [Example: Configuring MLPPP-over-ADSL Interfaces on page 124](#)
- [Example: Configuring the DHCP Client on ADSL Interface on page 126](#)
- [Example: Configuring CHAP on DSL Interfaces on page 130](#)
- [Example: Configuring ATM-over-ADSL Network Interfaces on page 138](#)

ADSL Interface Overview

Selected Juniper Networks security devices support DSL features including ATM-over-ADSL and ATM-over-SHDSL interfaces.



NOTE: Payload loopback functionality is not supported on ATM-over-SHDSL interfaces.

Asymmetric digital subscriber line (ADSL) technology is part of the xDSL family of modem technologies that use existing twisted-pair telephone lines to transport high-bandwidth data. ADSL lines connect service provider networks and customer sites over the "last mile" of the network—the loop between the service provider and the customer site.

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 26 on page 111](#).

Table 26: 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 26: Standard Bandwidths of DSL Operating Modes (continued)

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

ADSL, ADSL2, and ADSL2+ support the following standards:

- For Annex A:
 - ITU G.992.1 (ADSL)
- For Annex A only:
 - ANSI T1.413 Issue II
 - ITU G.992.3 (ADSL2)
 - ITU G.992.5 (ADSL2+)
- For Annex M:
 - ITU G.992.3 (ADSL2)
 - ITU G.992.5 (ADSL2+)
- For Annex B:
 - ITU G.992.1 (ADSL)
 - ITU G.992.3 (ADSL2)
 - ITU G.992.5 (ADSL2+)
- For Annex B only
 - ETSI TS 101 388 V1.3

The ADSL Mini-PIM facilitates a maximum of 10 virtual circuits on supported security devices.

Supported security devices with Mini-PIMs can use PPP over Ethernet over ATM (PPPoEoA) and PPP over ATM (PPPoA) to connect through ADSL lines only.

ADSL Systems

ADSL links run across twisted-pair telephone wires. When ADSL modems are connected to each end of a telephone wire, a dual-purpose ADSL circuit can be created. Once established, the circuit can transmit lower-frequency voice traffic and higher-frequency data traffic.

To accommodate both types of traffic, ADSL modems are connected to plain old telephone service (POTS) splitters that filter out the lower-bandwidth voice traffic and the higher-bandwidth data traffic. The voice traffic can be directed as normal telephone voice traffic. The data traffic is directed to the ADSL modem, which is typically connected to the data network.

ADSL2 and ADSL2+

The ADSL2 and ADSL2+ standards were adopted by the ITU in July 2002. ADSL2 improves the data rate and reach performance, diagnostics, standby mode, and interoperability of ADSL modems.

ADSL2+ doubles the possible downstream data bandwidth, enabling rates of 20 Mbps on telephone lines shorter than 5000 feet (1.5 km).

ADSL2 uses seamless rate adaptation (SRA) to change the data rate of a connection during operation with no interruptions or bit errors. The ADSL2 transceiver detects changes in channel conditions—for example, the failure of another transceiver in a multicarrier link—and sends a message to the transmitter to initiate a data rate change. The message includes data transmission parameters such as the number of bits modulated and the power on each channel. When the transmitter receives the information, it transitions to the new transmission rate.

ATM CoS Support

Certain class-of-service (CoS) components for Asynchronous Transmission Mode (ATM) are provided to control data transfer, especially for time-sensitive voice packets. The ADSL Mini-PIM on the SRX210 device provides extended ATM CoS functionality to provide cells across the network. You can define bandwidth utilization, which consists of either a constant rate or a peak cell rate, with sustained cell rate and burst tolerance. By default, unspecified bit rate (UBR) is used because the bandwidth utilization is unlimited.

The following ATM traffic shaping features are supported:

Constant bit rate (CBR)	CBR is the service category for traffic with rigorous timing requirements like voice and certain types of video. CBR traffic needs a constant cell transmission rate throughout the duration of the connection.
Variable bit rate non-real-time (VBR-NRT)	VBR-NRT is intended for sources such as data transfer, which do not have strict time or delay requirements. VBR-NRT is suitable for packet data transfers.
Unspecified bit rate (UBR)	UBR is ATM's best-effort service, which does not provide any CoS guarantees. This is suitable for noncritical applications that can tolerate or quickly adjust to loss of cells.

The ability of a network to guarantee class of service depends on the way in which the source generates cells and also on the availability of network resources. The connection contract between the user and the network thus contains information about the way in which traffic is generated by the source.

A set of traffic descriptors is specified for this purpose. The network provides the class of service for the cells that do not violate these specifications. The following are the traffic descriptors specified for an ATM network:

- 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)—The maximum burst size that can be sent at the peak rate.
- Cell delay variation tolerance (CDVT)—Allows the user to delay the traffic for a particular time duration in microseconds to follow a rhythmic pattern.

For traffic that does not require the ability to periodically burst to a higher rate, you can specify a CBR. You can configure VBR-NRT for ATM interfaces, which supports VBR data traffic with average and peak traffic parameters. VBR-NRT is scheduled with a lower priority and with a larger sustained cell rate (SCR) limit, allowing it to recover bandwidth if it falls behind.

On SRX300, SRX320, SRX340, SRX345, and SRX550HM 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.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 389](#)
- [ADSL and SHDSL Interfaces Configuration Overview on page 114](#)
- [Example: Configuring ATM-over-ADSL Network Interfaces on page 138](#)
- [Example: Configuring ATM-over-SHDSL Network Interfaces on page 118](#)
- [Example: Configuring CHAP on DSL Interfaces on page 130](#)
- [Example: Configuring MLPPP-over-ADSL Interfaces on page 124](#)

ADSL and SHDSL Interfaces Configuration Overview

An SRX Series device with an ADSL interface supports LFI through an MLPPP.



NOTE: Currently, Junos OS supports bundling of only one xDSL link under bundle interface.

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.

You can configure the following properties for the ADSL and SHDSL interfaces:

- Physical properties
- Logical properties

You can configure the following physical properties for the interface:

- ATM virtual path identifier (VPI) options for the interface—for example, at-2/0/0:
 - ATM VPI—A number from 0 through 255—for example, 25.
 - Operation, Maintenance, and Administration (OAM) F5 loopback cell thresholds (“liveness”) on ATM virtual circuits. The range is from 1 through 255, and the default is 5 cells.
 - Down count—Number of consecutive OAM loopback cells an ATM virtual circuit must lose to be identified as unavailable—for example, 200.
 - Up count—Number of consecutive OAM loopback cells an ATM virtual interface must receive to be identified as operational—for example, 200.
 - OAM period—Interval, in seconds, at which OAM cells are transmitted on ATM virtual circuits—for example, 100. The range is from 1 through 900 seconds.
- Configure CBR for the interface—for example, at-1/0/0.
 - CBR—Range from 33,000 through 1,199,920
 - CDVT—Range from 1 through 9,999
- Configure VBR for the interface—for example, at-1/0/0.
 - MBS—Range from 33,000 through 1,199,920
 - CDVT—Range from 1 through 9,999
 - PCR—Range from 33,000 through 1,199,920
 - SCR—Range from 33,000 through 1,199,920
- Type of DSL operating mode for the ATM-over-ADSL and ATM-over-SHDSL interfaces—for example, auto:

Annex A (used in North American network implementations) and Annex B (used in European network implementations) support the following operating modes:

- **auto**—Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex A, the ADSL interface trains in either ANSI T1.413 Issue II mode or ITU G.992.1 mode. For Annex B, the ADSL interface trains in ITU G.992.1 mode. For the SHDSL interface, the line rate is available only in two-wire mode and is the default value.
- **itu-dmt**—Configures the ADSL interface to train in ITU G.992.1 mode.
- **192 Kbps or higher**—Speed of transmission of data on the SHDSL connection. For the SHDSL interface, in the four-wire mode, the default line rate is 4,608 Kbps.

Annex A supports the following operating modes:

- **adsl2plus**—Configures the ADSL interface to train in ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM.
- **itu-dmt-bis**—Configures the ADSL interface to train in ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM.
- **ansi-dmt**—Configures the ADSL interface to train in the ANSI T1.413 Issue II mode.

Annex B supports the following operating modes:

- **etsi**—Configures the ADSL line to train in the ETSI TS 101 388 V1.3.1 mode.
- **itu-annexb-ur2**—Configures the ADSL line to train in the G.992.1 Deutsche Telekom UR-2 mode.
- **itu-annexb-non-ur2**—Configures the ADSL line to train in the G.992.1 Non-UR-2 mode.
- Loopback option for testing the SHDSL connection integrity—for example, local loopback.

The following values are available:

- **local**—Used for testing the SHDSL equipment with local network devices.
- **payload**—Used to command the remote configuration to send back the received payload.
- **remote**—Used to test SHDSL with a remote network configuration.
- Signal-to-noise ratio (SNR) margin—for example, 5 dB for either or both of the following thresholds:
 - **current**—Line trains at higher than current noise margin plus SNR threshold. The range is from 0 to 10 dB. The default value is 0.
 - **snext**—Line trains at higher than self-near-end crosstalk (SNEXT) threshold. The default value is **disabled**.

Setting the SNR creates a more stable SHDSL connection by making the line train at a SNR margin higher than the threshold. If any external noise below the threshold is applied to the line, the line remains stable. You can also disable the SNR margin thresholds.

- Encapsulation type—for example, ethernet-over-atm:
 - **atm-pvc**—ATM permanent virtual circuits is the default encapsulation for ATM-over-ADSL and ATM-over-SHDSL interfaces.

For PPP over ATM (PPPoA)-over-ADSL and over-SHDSL interfaces, use this type of encapsulation.
 - **ethernet-over-atm**—Ethernet over ATM encapsulation.

For PPP over Ethernet (PPPoE) over ATM-over-ADSL and ATM-over-SHDSL interfaces that carry IPv4 traffic, use this type of encapsulation.

You can configure the following logical properties for the interface:

- Logical interface. Set a value from 0 through 16,385—for example, 3. Add other values if required by your network.
- Configure encapsulation for the ATM-for-ADSL or ATM-for-SHDSL logical unit—for example, atm-nlpid.

The following encapsulations are supported on the ATM-over-ADSL and ATM-over-SHDSL interfaces that use inet (IP) protocols only:

- **atm-vc-mux**—Use ATM virtual circuit multiplex encapsulation.
- **atm-nlpid**—Use ATM network layer protocol identifier (NLPID) encapsulation.
- **atm-cisco-nlpid**—Use Cisco NLPID encapsulation.
- **ether-over-atm-llc**—For interfaces that carry IPv4 traffic, use Ethernet over LLC encapsulation. You cannot configure multipoint interfaces if you use this type of encapsulation.

The following encapsulations are supported on the ATM-over-ADSL or ATM-over-SHDSL for PPP-over-ATM (PPPoA) interfaces only.

- **atm-ppp-llc**—AAL5 logical link control (LLC) encapsulation.
- **atm-ppp-vc-mux**—Use AAL5 multiplex encapsulation.

Other encapsulation types supported on the ATM-over-ADSL and ATM-over-SHDSL interfaces are:

- **ppp-over-ether-over-atm-llc**—Use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead you configure the interface address on the PPP interface.
- **atm-snap**—Use ATM subnetwork attachment point (SNAP) encapsulation.
- OAM options for the ATM virtual circuits:
 - OAM F5 loopback cell thresholds (“liveness”) on ATM virtual circuits. The range is from 1 through 255, and the default is 5 cells.
 - Down count—Number of consecutive OAM loopback cells an ATM virtual circuit must lose to be identified as unavailable—for example, 200.
 - Up count—Number of consecutive OAM loopback cells an ATM virtual interface must receive to be identified as operational—for example, 200.
 - OAM period—Interval, in seconds, at which OAM cells are transmitted on ATM virtual circuits—for example, 100. The range is from 1 through 900 seconds.
- Family protocol type—for example, inet. Commands vary depending on the protocol type.
- ATM VCI options for the interface:
 - ATM VCI type—vci
 - ATM VCI value—A number from 0 through 4,089—for example, 35—with VCIs 0 through 31 reserved.

**Related
Documentation**

- [Understanding Point-to-Point Protocol over Ethernet on page 389.](#)
- [ADSL Interface Overview on page 111](#)

- [Example: Configuring ATM-over-ADSL Network Interfaces on page 138](#)
- [Example: Configuring ATM-over-SHDSL Network Interfaces on page 118](#)
- [Example: Configuring CHAP on DSL Interfaces on page 130](#)
- [Example: Configuring MLPPP-over-ADSL Interfaces on page 124](#)

Example: Configuring ATM-over-SHDSL Network Interfaces

This example shows how to configure ATM-over-SHDSL network interfaces.

- [Requirements on page 118](#)
- [Overview on page 118](#)
- [Configuration on page 118](#)
- [Verification on page 121](#)

Requirements

Before you begin:

- Configure network interfaces as necessary. See [“Understanding Ethernet Interfaces” on page 259](#).
- Configure PPPoE encapsulation on an Ethernet interface or on an ATM-over-ADSL interface. See [“Understanding Point-to-Point Protocol over Ethernet” on page 389](#).

Overview

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.

Configuration

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 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 snxt 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

```

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 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
```

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 25 {
    oam-period 100;
    oam-liveness {
      up-count 200;
      down-count 200;
    }
  }
}
shdsl-options {
```



```

annex annex-a;
line-rate auto;
loopback local;
snr-margin {
    current 5
    snr 5;
}
}
unit 3 {
    encapsulation atm-nlpid;
    vci 35;
    oam-period 100;
    oam-liveness {
        up-count 200;
        down-count 200;
    }
    family inet;
}

```

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

Verification

Confirm that the configuration is working properly.

Verifying an ATM-over-SHDSL Configuration

Purpose Verify that the interface properties are correct.

Action From operational mode, enter the **show interfaces at-2/0/0 extensive** command.

```

user@host> show interfaces at-2/0/0 extensive
Physical interface: at-2/0/0, 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

```

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 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 page (*Interfaces*>*interface-name*).
- 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 expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics interface-name** command.
- No SHDSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm.
 - **LOS**—Loss of signal. No signal was detected on the line.
 - **LOSW**—Loss of sync word. A message ID was sent.
 - **Power status**—A power failure has occurred.
 - **LOSD**—Loss of signal was detected at the remote application interface.
 - **ES**—Errored seconds. One or more cyclic redundancy check (CRC) anomalies were detected.
 - **SES**—Severely errored seconds. At least 50 CRC anomalies were detected.
 - **UAS**—Unavailable seconds. An interval has occurred during which one or more LOSW defects were detected.

Examine the SHDSL interface status:

- **Line termination**—SHDSL transceiver unit—remote (STU—R). (Only customer premises equipment is supported.)
- **Annex**—Either Annex A or Annex B. Annex A is supported in North America, and Annex B is supported in Europe.
- **Line mode**—SHDSL mode configured on the G.SHDSL interface pair, either two-wire or four-wire.
- **Modem status**—Data. Sending or receiving data.
- **Last fail code**—Code for the last interface failure.
- **Framer mode**—ATM Framer mode of the underlying interface.
- **Chipset version**—Version number of the chipset on the interface
- **Firmware version**—Version number of the firmware on the interface.

Examine the operational statistics for a SHDSL interface.

- **Loop attenuation (dB)**—Reduction in signal strength measured in decibels.
- **Transmit power (dB)**—Amount of SHDSL usage in %.
- **Receiver gain (dB)**—Maximum extraneous signal allowed without causing the output to deviate from an acceptable level.

- **SNR sampling (dB)**—Signal-to-noise ratio at a receiver point in decibels.
- **Bit rate (kbps)**—Data transfer speed on the SHDSL interface.
- **CRC errors**—Number of cyclic redundancy check errors.
- **SEGA errors**—Number of segment anomaly errors. A regenerator operating on a segment received corrupted data.
- **LOSW errors**—Number of loss of signal defect errors. Three or more consecutively received frames contained one or more errors in the framing bits.
- **Received cells**—Number of cells received through the interface.
- **Transmitted cells**—Number of cells sent through the interface.
- **HEC errors**—Number of header error checksum errors.
- **Cell drop**—Number of dropped cells on the interface.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [ADSL Interface Overview on page 111](#)
- [ADSL and SHDSL Interfaces Configuration Overview on page 114](#)
- [Example: Configuring CHAP on DSL Interfaces on page 130](#)

Example: Configuring MLPPP-over-ADSL Interfaces

This example shows how to configure MLPPP on an ADSL interface.

- [Requirements on page 124](#)
- [Overview on page 124](#)
- [Configuration on page 125](#)
- [Verification on page 125](#)

Requirements

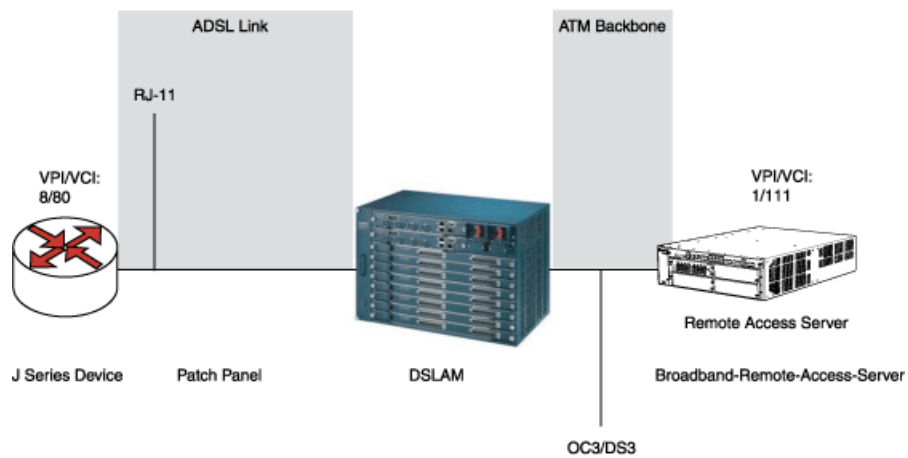
Before you begin, configure network interfaces as necessary. See “[Understanding Ethernet Interfaces](#)” on page 259.

Overview

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 10 on page 125](#) shows a typical example of MLPPP-over-ADSL end-to-end connectivity.

Figure 10: MLPPP-over-ADSL Interface



Configuration

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 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
```

Verification

To verify the configuration is working properly, enter the **show interfaces at-5/0/0** command.

Related Documentation

- [Understanding Interfaces on page 3](#)

- [ADSL Interface Overview on page 111](#)
- [ADSL and SHDSL Interfaces Configuration Overview on page 114](#)

Example: Configuring the DHCP Client on ADSL Interface

This example shows how to configure DHCP client on ADSL or SHDSL or VDSL2 interface (when VDSL2 interface is configured to operate in ADSL fallback mode).

- [Requirements on page 126](#)
- [Overview on page 126](#)
- [Configuration on page 126](#)
- [Verification on page 129](#)

Requirements

Before you begin:

- Review the overview section on DHCP client. See *Understanding DHCP Client Operation*
- Establish basic connectivity. See the Quick Start for your device.
- Configure network interfaces as necessary. See “[Example: Creating an Ethernet Interface](#)” on page 265.

Overview

In this example, you configure the ATM interface as **at-1/0/0**. You then set the logical interface to unit 0 and specify the family protocol type as inet. Finally, you configure the DHCP client.

Configuration

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-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
```

Step-by-Step Procedure

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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-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 at-1/0/0
encapsulation ethernet-over-atm;
atm-options {
  vpi 2;
}
dsl-options {
  operating-mode auto;
}
unit 0 {
  encapsulation ether-over-atm-llc;
  vci 2.122;
  family inet {
    dhcp {
      client-identifier ascii 00:0a:12:00:12:12;
      lease-time 86400;
      retransmission-attempt 6;
      retransmission-interval 5;
      server-address 10.1.1.1;
    }
  }
}
```


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

Verification

Confirm that the configuration is working properly.

- [Verifying the DHCP Configuration on page 129](#)
- [Verify Interface Status on page 129](#)

Verifying the DHCP Configuration

Purpose Verify that the DHCP options are configured properly.

Action Verify the DHCP configuration by using 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
```

Verify Interface Status

Purpose Verify the interface status and check traffic statistics.

Action Verify interface status by using 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
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 10.40.1.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 20.086/26.404/61.723/6.194 ms
```

- Related Documentation**
- [DHCP Server Configuration Overview](#)
 -

Example: Configuring CHAP on DSL Interfaces

This example shows how to configure CHAP on either the ATM-over-ADSL or the ATM-over-SHDSL interface.

- [Requirements on page 130](#)
- [Overview on page 130](#)
- [Configuration on page 130](#)
- [Verification on page 131](#)

Requirements

Before you begin, configure network interfaces as necessary. See “[Understanding Ethernet Interfaces](#)” on page 259.

Overview

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.

Configuration

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 at-3/0/0 unit 0 ppp-options chap access-profile A-ppp-client local-name
A-at-3/0/0.0 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 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
```

Results From configuration mode, confirm your configuration by entering the **show access profile A-ppp-client** and **show interfaces at-3/0/0** 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 A-ppp-client
client client1 chap-secret "$9$ikPQtu1Sre0BclMW-dk.P5QnApB"; ## SECRET-DATA
[edit]
user@host# show interfaces at-3/0/0
unit 0 {
  ppp-options {
    chap {
      access-profile A-ppp-client;
      local-name A-at-3/0/0.0;
      passive;
    }
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying ADSL Interface Properties on page 132](#)
- [Verifying a PPPoA Configuration for an ATM-over-ADSL Interface on page 134](#)
- [Verifying an ATM-over-SHDSL Configuration on page 135](#)

Verifying ADSL Interface Properties

Purpose Verify that the ADSL interface properties are enabled.

Action From operational mode, enter 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:
    Vendor Country : 0x0f 0xb5
    Vendor ID      : STMI IFTN
    Vendor Specific: 0x0000 0x70de
  ADSL Statistics:
    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: 0, Input 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 minimipim does not send any alarms. So we can't show alarm stats for minimipim. So following information will not be displayed in Minimipim 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:
  Vendor Country :          0x0f          0xb5
  Vendor ID      :          STMI          IFTN
  Vendor Specific:          0x0000          0x70de

```

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 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 page (Interfaces>interface-name).
- The physical link is up. A link state of dDown 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 inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics interface-name** command.
- No ADSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm. The following are ADSL-specific alarms:
 - **LOCDI**—Loss of cell delineation for interleaved channel
 - **LOCDNI**—Loss of cell delineation for noninterleaved channel
 - **LOF**—Loss of frame
 - **LOM**—Loss of multiframe
 - **LOP**—Loss of power
 - **LOS**—Loss of signal

Examine the operational statistics for an ADSL interface. Statistics in the ATU-R (ADSL transceiver unit—remote) column are for the near end. Statistics in the ATU-C (ADSL transceiver unit—central office) column are for the far end.

- **Attenuation (dB)**—Reduction in signal strength measured in decibels.
- **Capacity used (%)**—Amount of ADSL usage in %.
- **Noise margin (dB)**—Maximum extraneous signal allowed without causing the output to deviate from an acceptable level.
- **Output power (dBm)**—Amount of power used by the ADSL interface.
- **Bit rate (kbps)**—Data transfer speed on the ADSL interface.

Verifying a PPPoA Configuration for an ATM-over-ADSL Interface

Purpose Verify that the PPPoA configuration for an ATM-over-ADSL interface is correct.

Action From operational mode, enter the **show interfaces at-3/0/0** and the **show access** commands.

Verifying an ATM-over-SHDSL Configuration

Purpose Verify that the interface properties are correct.

Action From operational mode, enter 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
```

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 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 page (Interfaces>*interface-name*).
- 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 inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics *interface-name*** command.
- No SHDSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm.
 - **LOS**—Loss of signal. No signal was detected on the line.
 - **LOSW**—Loss of sync word. A message ID was sent.
 - **Power status**—A power failure has occurred.
 - **LOSD**—Loss of signal was detected at the remote application interface.
 - **ES**—Errored seconds. One or more cyclic redundancy check (CRC) anomalies were detected.

- **SES**—Severely errored seconds. At least 50 CRC anomalies were detected.
- **UAS**—Unavailable seconds. An interval has occurred during which one or more LOSW defects were detected.

Examine the SHDSL interface status:

- **Line termination**—SHDSL transceiver unit—remote (STU—R). (Only customer premises equipment is supported.)
- **Annex**—Either Annex A or Annex B. Annex A is supported in North America, and Annex B is supported in Europe.
- **Line mode**—SHDSL mode configured on the G.SHDSL interface pair, either two-wire or four-wire.
- **Modem Status**—Data. Sending or receiving data.
- **Last fail code**—Code for the last interface failure.
- **Framer mode**—Framer mode of the underlying interface: ATM.
- **Dying gasp**—Ability of a device that has lost power to send a message informing the attached DSL access multiplexer (DSLAM) that it is about to go offline.
- **Chipset version**—Version number of the chipset on the interface
- **Firmware version**—Version number of the firmware on the interface.

Examine the operational statistics for a SHDSL interface.

- **Loop attenuation (dB)**—Reduction in signal strength measured in decibels.
- **Transmit power (dB)**—Amount of SHDSL usage in %.
- **Receiver gain (dB)**—Maximum extraneous signal allowed without causing the output to deviate from an acceptable level.
- **SNR sampling (dB)**—Signal-to-noise ratio at a receiver point in decibels.
- **Bit rate (kbps)**—Data transfer speed on the SHDSL interface.
- **CRC errors**—Number of cyclic redundancy check errors.
- **SEGA errors**—Number of segment anomaly errors. A regenerator operating on a segment received corrupted data.
- **LOSW errors**—Number of loss of signal defect errors. Three or more consecutively received frames contained one or more errors in the framing bits.
- **Received cells**—Number of cells received through the interface.
- **Transmitted cells**—Number of cells sent through the interface.
- **HEC errors**—Number of header error checksum errors.
- **Cell drop**—Number of dropped cells on the interface.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [ADSL Interface Overview on page 111](#)
- [ADSL and SHDSL Interfaces Configuration Overview on page 114](#)
- [Example: Configuring MLPPP-over-ADSL Interfaces on page 124](#)

Example: Configuring ATM-over-ADSL Network Interfaces

This example shows how to configure ATM-over-ADSL network interfaces for the devices.

- [Requirements on page 138](#)
- [Overview on page 138](#)
- [Configuration on page 139](#)
- [Verification on page 141](#)

Requirements

Before you begin:

- Configure network interfaces as necessary. See [“Understanding Ethernet Interfaces” on page 259](#).
- Configure PPPoE encapsulation on an Ethernet interface or on an ATM-over-ADSL interface. See [“Understanding Point-to-Point Protocol over Ethernet” on page 389](#).

Overview

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.



NOTE: On SRX300, SRX320, SRX340, SRX345, and SRX550HM 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.

Configuration

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 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
```

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 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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show interfaces at-2/0/0** 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 at-1/0/0
encapsulation ethernet-over-atm;
dsl-options {
  operating-mode auto;
}
unit 0 {
  shaping {
    vbr peak 33k;
    burst
  }
}
unit 3 {
  encapsulation atm-nlpid;
```

```

vci 35;
oam-period 100;
oam-liveness {
    up-count 200;
    down-count 200;
}
family inet;
}
[edit]
user@host show interfaces at-2/0/0
atm-options {
vpi 25 {
    oam-period 100;
    oam-liveness {
up-count 200;
down-count 200
    }
}
}
}

```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the ADSL Interface Properties on page 141](#)
- [Verifying a PPPoA Configuration for an ATM-over-ADSL Interface on page 144](#)

Verifying the ADSL Interface Properties

Purpose Verify that the interface properties are correct.

Action From operational mode, enter 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:
  Vendor Country : ATU-R 0x0f ATU-C 0xb5
  Vendor ID      : STMI  IFTN
  Vendor Specific: 0x0000 0x70de
ADSL Statistics:
  Attenuation (dB) : ATU-R 0.0 ATU-C 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: 0, Input 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
Limit      Bandwidth      Buffer Priority
           %      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 minimipim does not send any alarms. So we can't show alarm stats for minimipim. So following information will not be displayed in Minimipim 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:
Vendor Country :      0x0f      0xb5
Vendor ID      :      STMI      IFTN
Vendor Specific:    0x0000    0x70de

```

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 either of the following:
 - In the CLI, delete the **disable** statement at the [edit interfaces *interface-name*] level of the configuration hierarchy.
 - In J-Web, clear the **Disable** check box on the Interfaces page (Interfaces>*interface-name*).
- 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 inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics *interface-name*** command.
- No ADSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm. The following are ADSL-specific alarms:
 - **LOCDI**—Loss of cell delineation for interleaved channel.
 - **LOCDNI**—Loss of cell delineation for noninterleaved channel.
 - **LOF**—Loss of frame.
 - **LOM**—Loss of multiframe.
 - **LOP**—Loss of power.
 - **LOS**—Loss of signal.

Examine the operational statistics for an ADSL interface. Statistics in the ATU-R (ADSL transceiver unit—remote) column are for the near end. Statistics in the ATU-C (ADSL transceiver unit—central office) column are for the far end.

- **Attenuation (dB)**—Reduction in signal strength .
- **Capacity used (%)**—Amount of ADSL usage.
- **Noise margin (dB)**—Maximum extraneous signal allowed without causing the output to deviate from an acceptable level.
- **Output power (dBm)**—Amount of power used by the ADSL interface.
- **Bit rate (kbps)**—Data transfer speed on the ADSL interface.

Verifying a PPPoA Configuration for an ATM-over-ADSL Interface

Purpose Verify that the PPPoA configuration for an ATM-over-ADSL interface is correct.

Action From operational mode, enter the **show interfaces at-1/0/0** and the **show access** commands.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [ADSL Interface Overview on page 111](#)
- [ADSL and SHDSL Interfaces Configuration Overview on page 114](#)
- [Example: Configuring ATM-over-SHDSL Network Interfaces on page 118](#)
- [Example: Configuring MLPPP-over-ADSL Interfaces on page 124](#)

Configuring G.SHDSL Interfaces

- [SHDSL Interface Overview on page 145](#)
- [G.SHDSL Mini-PIM Overview on page 146](#)
- [G.SHDSL Mini-PIM Configuration Overview on page 148](#)
- [Example: Configuring the G.SHDSL Interface on page 150](#)
- [Example: Configuring the G.SHDSL Interface on SRX Series Devices on page 158](#)
- [Example: Configuring the G.SHDSL Interface in EFM Mode on page 170](#)

SHDSL Interface Overview

Symmetric high-speed DSL (SHDSL) interfaces on some SRX Series devices support 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.

Unlike ADSL, which delivers more bandwidth downstream than available upstream, SHDSL is symmetrical and delivers a bandwidth of up to 2.3 Mbps in both directions. Because business applications require high-speed digital transportation methods, SHDSL is becoming very popular and gaining wide acceptance in the industry. Additionally, SHDSL is compatible with ADSL and therefore causes very little, if any, interference between cables.

SHDSL is deployed on a network in much the same manner as ADSL.

SHDSL interfaces support Packet Transfer Mode (PTM). In PTM, 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.



NOTE: Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Related Documentation

- [G.SHDSL Mini-PIM Overview on page 146](#)
- [G.SHDSL Mini-PIM Configuration Overview on page 148](#)
- [Example: Configuring the G.SHDSL Interface on page 150](#)
- [Example: Configuring the G.SHDSL Interface on SRX Series Devices on page 158](#)
- [Example: Configuring the G.SHDSL Interface in EFM Mode on page 170](#)

G.SHDSL Mini-PIM Overview

Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

The G.SHDSL Mini-Physical Interface Module (Mini-PIM) provides the physical connection to DSL network media types.

The G.SHDSL Mini-PIM provides the following Asynchronous Transfer Mode (ATM) key features:

- 2-wire (4-port 2-wire) mode, 4-wire (2-port 4-wire) mode, and 8-wire (1-port 8-wire) mode support
- Virtual circuits (VC) per Mini-PIM (10 maximum including OAM VC)
- ATM-over-G.SHDSL framing
- ATM OAM support
- Maximum MTU size of 9180 bytes
- Noise margin support
- Point-to-Point Protocol over ATM and PPPoE over ATM encapsulation support
- Local loopback mode support
- Dying gasp support

The G.SHDSL Mini-PIM provides extended ATM CoS functionality to cells across the network. You can define bandwidth utilization, which consists of either a constant rate or a peak cell rate, with sustained cell rate and burst tolerance. By default, unspecified bit rate (UBR) is used because the bandwidth utilization is unlimited.

The following ATM traffic shaping features are supported:

- **Constant bit rate (CBR)**—CBR is the service category for traffic with rigorous timing requirements like voice and certain types of video. CBR traffic needs a constant cell transmission rate throughout the duration of the connection.
- **Variable bit rate, non-real-time (VBR-NRT)**—VBR-NRT is intended for sources such as data transfer, which do not have strict time or delay requirements. VBR-NRT is suitable for packet data transfers.
- **Variable bit rate, real-time (VBR-RT)**—VBR-RT is intended for sources such as data transfer, which takes place in real time. VBR-RT requires access to time slots at a rate that can vary significantly from time to time.

Table 27 on page 147 displays the traffic descriptors specified for an ATM network.

Table 27: Traffic Descriptors

Traffic Descriptors	Description
Peak cell rate (PCR)	Maximum 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.

The G.SHDSL Mini-PIM provides the following Packet Transfer Mode (PTM) Ethernet in the First Mile (EFM) key features:

- EFM PIC mode support
- Maximum MTU size of 1514 bytes
- PPPoE encapsulation support
- Local loopback mode support
- Chassis cluster mode support
- Dying gasp support
- IPv6 support
- VLAN over EFM support

The following four annexes are supported on the G.SHDSL Mini-PIM in both ATM and PTM EFM modes:

- Annex A
- Annex B
- Annex F
- Annex G

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 SRX210, SRX220, SRX240, and SRX550 devices using the symmetrical WAN speeds shown in [Table 28 on page 148](#).

Table 28: 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.		

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [SHDSL Interface Overview on page 145](#)
- [G.SHDSL Mini-PIM Configuration Overview on page 148](#)
- [Example: Configuring the G.SHDSL Interface on page 150](#)
- [Example: Configuring the G.SHDSL Interface on SRX Series Devices on page 158](#)
- [Example: Configuring the G.SHDSL Interface in EFM Mode on page 170](#)

G.SHDSL Mini-PIM Configuration Overview



NOTE: Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Specify the wire mode on the G.SHDSL interface using one of the following options:

- **1-port-atm**—Configures an 8-wire (1-port, 8-wire) wire mode.
- **2-port-atm**—Configures a 4-wire (2-port, 4-wire) wire mode.
- **4-port-atm**—Configures a 2-wire (4-port, 2-wire) wire mode.
- **efm**—Configures an efm (1-port, 2-wire) wire mode.



NOTE: The default wire mode is 4-wire (2-port, 4-wire).

Specify the annex type using one of the following options:

- Annex A
- Annex B
- Annex F
- Annex G



NOTE: The default annex type is auto.

Specify the SHDSL line rate (speed of transmission of data on the SHDSL connection) using one of the following values:

- **auto**—Automatically selects a line rate.
- **value**—Selects a value between 192 kbps and 22,784 kbps.



NOTE: The default line rate is auto.

Specify the encapsulation type using one of the following values:



NOTE: The pt interface does not require encapsulation types.

The at interface encapsulation types are as follows:

- **atm-pvc**—ATM permanent virtual circuits is the default encapsulation for ATM-over-SHDSL interfaces. For PPP over ATM (PPPoA) over SHDSL interfaces, use this type of encapsulation. Use this type of encapsulation if you are using ATM DSLAM.
- **ethernet-over-atm**—Ethernet over ATM encapsulation. For PPP over Ethernet (PPPoE) over ATM-over-SHDSL interfaces that carry IPv4 traffic, use this type of encapsulation. Use this type of encapsulation if you are using IP DSLAM.

Configure the encapsulation type using one of the following values:

- **atm-cisco-nlpid**—Cisco NLPID encapsulation.
- **atm-mlppp-llc**—ATM MLPPP over AAL5/LLC encapsulation.
- **atm-nlpid**—ATM Network Layer protocol identifier (NLPID) encapsulation.
- **atm-ppp-llc**—AAL5 logical link control (LLC) encapsulation.
- **atm-ppp-vc-mux**—AAL5 multiplex encapsulation.
- **atm-vc-mux**—ATM virtual circuit multiplex encapsulation.
- **atm-snap**—ATM subnetwork attachment point (SNAP) encapsulation.
- **ether-over-atm-llc**—For interfaces that carry IPv4 traffic, use Ethernet over LLC encapsulation. You cannot configure multipoint interfaces if you use this type of encapsulation.
- **ppp-over-ether-over-atm-llc**—PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead you configure the interface address on the PPP interface.

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [SHDSL Interface Overview on page 145](#)
- [G.SHDSL Mini-PIM Overview on page 146](#)
- [Example: Configuring the G.SHDSL Interface on page 150](#)
- [Example: Configuring the G.SHDSL Interface on SRX Series Devices on page 158](#)
- [Example: Configuring the G.SHDSL Interface in EFM Mode on page 170](#)

Example: Configuring the G.SHDSL Interface

This example shows how to configure the G.SHDSL interface.



NOTE: Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

- [Requirements on page 151](#)
- [Overview on page 151](#)

- [Configuration on page 151](#)
- [Verification on page 152](#)

Requirements

Before you begin, configure network interfaces as necessary. See [“Understanding Ethernet Interfaces” on page 259](#).

Overview

In this example, you specify the wire mode called 2-port-atm and create an interface called at-1/0/0. You then specify the annex type as annex-a and set the line rate to auto. Then you specify the encapsulation type as ethernet-over-atm and define a logical unit as unit 3 that you connect to this physical G.SHDSL interface. You can set a value from 0 through 7. Finally, you configure the encapsulation type as ether-over-atm-llc.

Configuration

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 chassis fpc 1 pic 0 shdsl pic-mode 2-port-atm
set interfaces at-1/0/0 shdsl-options annex annex-a line-rate auto
set interfaces at-1/0/0 encapsulation ethernet-over-atm
set interfaces at-1/0/0 unit 3 encapsulation ether-over-atm-llc
```

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 G.SHDSL interface:

1. Specify the wire mode.

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



NOTE: For configuring the G.SHDSL interface in chassis cluster mode, provide the node id also. For example, to configure an shdsl 2 port pic-mode in chassis cluster mode for the fpc slot 1 on the node 0, use the following command:

```
set chassis node 0 fpc 1 pic 0 shdsl pic-mode 2-port-atm
```

2. Create an interface.

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

3. Specify the annex type.

```
[edit interfaces at-1/0/0 shdsl-options]  
user@host# set annex annex-a
```
4. Configure the line rate.

```
[edit interfaces at-1/0/0 shdsl-options]  
user@host# set line-rate auto
```
5. Specify the encapsulation type.

```
[edit interfaces at-1/0/0]  
user@host# set encapsulation ethernet-over-atm
```
6. Define one or more logical units.

```
[edit interfaces at-1/0/0]  
user@host# edit unit 3
```
7. Configure the encapsulation type.

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

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show chassis fpc 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 at-1/0/0  
encapsulation ethernet-over-atm;  
shdsl-options {  
  annex annex-a;  
  line-rate auto;  
}  
unit 3 {  
  encapsulation ether-over-atm-llc;  
}  
[edit]  
user@host# show chassis fpc 1  
pic 0 {  
  shdsl {  
    pic-mode 2-port-atm;  
  }  
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying G.SHDSL Interface Properties

Purpose Verify that the G.SHDSL interface properties are configured properly.

Action From operational mode, enter the **show interfaces at-1/0/0 extensive** command.

```
user@host> show interfaces at-1/0/0 extensive
```

Four-wire mode for interface at-1/0/0:

Physical interface: at-1/0/0, Enabled, Physical link is Up

Interface index: 146, SNMP ifIndex: 139, Generation: 329

Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, Speed: SHDSL(4-wire)

Speed: SHDSL(4-wire), 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: 00:24:dc:01:cf:a0

Last flapped : 2009-09-24 00:19:03 PDT (00:00:54 ago)

Statistics last cleared: 2009-09-24 00:18:24 PDT (00:01:33 ago)

Traffic statistics:

Input bytes :	125	0 bps
Output bytes :	96	0 bps
Input packets:	2	0 pps
Output packets:	1	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: 1, 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	1	1	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

SHDSL alarms : None

SHDSL defects : None

SHDSL media:	Seconds	Count	State
LINE1_LOSD	32	0	OK
LINE1_LOSW	37	0	OK
LINE2_LOSD	32	0	OK
LINE2_LOSW	37	0	OK
ES	37		
SES	37		
UAS	48		

SHDSL status:
 Line termination : STU-R
 Annex : Annex B
 Line mode : 4-wire
 Modem status : Data
 Bit rate (kbps) : 4608
 Last fail mode : No failure (0x00)
 Frammer mode : ATM
 Dying gasp : Enabled
 Frammer sync status : In sync
 Chipset version : 00

SHDSL statistics:
 Loop attenuation (dB) : 0.0
 Transmit power (dBm) : 0.0
 Receiver gain (dB) : -inf
 SNR sampling (dB) : inf
 CRC errors : 0
 SEGA errors : 0
 LOSW errors : 0
 Received cells : 0
 Transmitted cells : 0
 HEC errors : 0
 Cell drop : 0

Packet Forwarding Engine configuration:

Destination slot: 1

CoS information:

Direction : Output

CoS transmit queue	Bandwidth	Buffer Priority
Limit		
	% bps % usec	
0 best-effort	95 4377600 95 0	low
none		
3 network-control	5 230400 5 0	low
none		

Logical interface at-1/0/0.0 (Index 76) (SNMP ifIndex 133) (Generation 402)

Flags: Point-To-Multipoint SNMP-Traps 0x0 Encapsulation: Ether-over-ATM-LLC

Traffic statistics:

Input bytes : 125
 Output bytes : 116
 Input packets: 2
 Output packets: 1

Local statistics:

Input bytes : 125
 Output bytes : 116
 Input packets: 2
 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: 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: 1468, Generation: 322, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 17.1.1/24, Local: 17.1.1.1, Broadcast: 17.1.1.255, Generation:
496
    VCI 1.70
    Flags: Active, Multicast
    Total down time: 0 sec, Last down: Never
    ATM per-VC transmit statistics:
    Tail queue packet drops: 0
    Traffic statistics:
    Input bytes :                     0
    Output bytes :                     0
    Input packets:                     0
    Output packets:                     0

Logical interface at-1/0/0.32767 (Index 77) (SNMP ifIndex 141) (Generation 403)

    Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0 Encapsulation: ATM-VCMUX

    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
    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
VCI 1.4
Flags: Active
Total down time: 0 sec, Last down: Never
ATM per-VC transmit statistics:
Tail queue packet drops: 0
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
```

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 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 page (*Interfaces*>*interface-name*).
- 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 expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics interface-name** command.
- No SHDSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm.
 - **LOS**—Loss of signal. No signal was detected on the line.
 - **LOSW**—Loss of sync word. A message ID was sent.
 - **Power status**—A power failure has occurred.
 - **LOSD**—Loss of signal was detected at the remote application interface.
 - **ES**—Errored seconds. One or more cyclic redundancy check (CRC) anomalies were detected.
 - **SES**—Severely errored seconds. At least 50 CRC anomalies were detected.
 - **UAS**—Unavailable seconds. An interval has occurred during which one or more LOSW defects were detected.

Examine the SHDSL interface status:

- **Line termination**—SHDSL transceiver unit—remote (STU—R). (Only customer premises equipment is supported.)
- **Annex**—Either Annex A or Annex B. Annex A is supported in North America, and Annex B is supported in Europe.
- **Line mode**—SHDSL mode configured on the G.SHDSL interface pair, either two-wire or four-wire.
- **Modem status**—Data. Sending or receiving data.
- **Bit rate (kbps)**—Data transfer speed on the SHDSL interface.
- **Last fail code**—Code for the last interface failure.
- **Framer mode**—ATM framer mode of the underlying interface.
- **Dying gasp**—Ability of a device that has lost power to send a message informing the attached DSLAM that it is about to go offline.
- **Chipset version**—Version number of the chipset on the interface

Examine the operational statistics for a SHDSL interface.

- **Loop attenuation (dB)**—Reduction in signal strength.
- **Transmit power (dB)**—Amount of SHDSL.

- **Receiver gain (dB)**—Maximum extraneous signal allowed without causing the output to deviate from an acceptable level.
- **SNR sampling (dB)**—Signal-to-noise ratio at a receiver point.
- **CRC errors**—Number of cyclic redundancy check errors.
- **SEGA errors**—Number of segment anomaly errors. A regenerator operating on a segment received corrupted data.
- **LOSW errors**—Number of loss of signal defect errors. Three or more consecutively received frames contained one or more errors in the framing bits.
- **Received cells**—Number of cells received through the interface.
- **Transmitted cells**—Number of cells sent through the interface.
- **HEC errors**—Number of header error checksum errors.
- **Cell drop**—Number of dropped cells on the interface.

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [SHDSL Interface Overview on page 145](#)
- [G.SHDSL Mini-PIM Overview on page 146](#)
- [G.SHDSL Mini-PIM Configuration Overview on page 148](#)
- [Example: Configuring the G.SHDSL Interface on SRX Series Devices on page 158](#)

Example: Configuring the G.SHDSL Interface on SRX Series Devices

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



NOTE: Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

- [Requirements on page 159](#)
- [Overview on page 159](#)
- [Configuration on page 161](#)
- [Verification on page 169](#)

Requirements

Before you begin:

- Configure the network interfaces as necessary. See [“Understanding Ethernet Interfaces” on page 259](#).
- 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).



NOTE: This example uses an SRX210 Services Gateway. The information is also applicable to the SRX220 and SRX240 devices.

Overview

Figure 11 on page 159 shows the topology for the G.SHDSL Mini-PIM operating in 2X4-wire mode.

Figure 11: G.SHDSL Mini-PIM Operating in 2X4-Wire Mode




Figure 12 on page 159 shows the topology for the G.SHDSL Mini-PIM operating in 4X2-wire mode.

Figure 12: G.SHDSL Mini-PIM Operating in 4X2-Wire Mode





Figure 13 on page 160 shows the topology for the G.SHDSL Mini-PIM operating in 1X8-wire mode.

Figure 13: 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 29 on page 160.

Table 29: 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



NOTE: When the wire mode is set 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, you first configure a basic G.SHDSL interface. You set the operation wire mode to 2-port-atm, the line rate to 4096, and the annex type to annex-a.

You then configure the G.SHDSL interface when the device is connected to an IP DSLAM. You set the type of encapsulation to ethernet-over-atm and the ATM VPI option to 0. Then you 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, you set the interface address for the logical interface to 1.1.1.1/24.

Then you configure the G.SHDSL interface when the device is connected to an ATM DSLAM. You set the type of encapsulation to atm-pvc and the ATM VPI to 0. Then you set the type of encapsulation on the G.SHDSL logical interface to atm-snap and the ATM VCI to 0.65. Also, you set the interface address for the logical interface to 2.1.1.1/24

Next you configure PPPoE over ATM for the G.SHDSL Interface. You then set the ATM VPI to 0 and set the type of encapsulation to `ppp-over-ether-over-atm-llc`. 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 `at-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.) You then specify the logical interface as the client for the PPPoE interface and obtain an IP address by negotiation with the remote end.

Finally, you configure PPPoA over ATM for the G.SHDSL Interface. You set the type of encapsulation to `atm-pvc` and the ATM VPI to 0. You then 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. You 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.

Configuration

- [Configuring a Basic G.SHDSL Interface on page 161](#)
- [Configuring a G.SHDSL Interface When Connected to an IP DSLAM on page 162](#)
- [Configuring a G.SHDSL Interface When Connected to an ATM DSLAM on page 163](#)
- [Configuring PPPoE over ATM for the G.SHDSL Interface on page 165](#)
- [Configuring PPPoA over ATM for the G.SHDSL Interface on page 167](#)

Configuring a Basic G.SHDSL 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 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
```

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 quickly configure a basic G.SHDSL interface:

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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show chassis fpc 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 at-1/0/0
shdsl-options {
  annex annex-a;
  line-rate 4096;
}
[edit]
user@host# show chassis fpc 1
pic 0 {
  shdsl {
    pic-mode 2-port-atm;
  }
}
```

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

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

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-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 1.1.1.1/24
```

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 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 1.1.1.1/24
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-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 at-1/0/0
encapsulation ethernet-over-atm;
atm-options {
  vpi 0;
}
unit 0 {
  encapsulation ether-over-atm-llc;
  vci 0.60;
  family inet {
    address 1.1.1.1/24;
  }
}
```

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

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

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-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 2.1.1.1/24
```

**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 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 2.1.1.1/24
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-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 at-1/0/0
encapsulation atm-pvc;
```

```

atm-options {
  vpi 0;
}
unit 0 {
  encapsulation atm-snap;
  vci 0.65;
  family inet {
    address 2.1.1/24
  }
}

```

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

Configuring PPPoE over ATM for the G.SHDSL 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 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

```

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 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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show interfaces pp0** 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 at-1/0/0
encapsulation ethernet-over-atm;
atm-options {
  vpi 0;
}
unit 0 {
  encapsulation ppp-over-ether-over-atm-llc;
  vci 0.35;
}
[edit]
user@host# show interfaces pp0
unit 0 {
  ppp-options {
    pap {
      access-profile pap_prof;
    }
    local-name srx-210;
    local-password "$9$0tLw1SeN-woJDSr-wY2GU69Cp1RSre";
    passive;
  }
  pppoe-options {
    underlying-interface at-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    negotiate-address;
  }
}
```

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

Configuring PPPoA over ATM for the G.SHDSL 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 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
```

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 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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-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 at-1/0/0
encapsulation atm-pvc;
atm-options {
  vpi 0;
}
unit 0 {
  encapsulation atm-ppp-llc;
  vci 1.122;
  ppp-options {
    chap {
      access-profile juniper;
      local-name srx-210;
    }
  }
  family inet {
    negotiate-address;
  }
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying G.SHDSL Interface Properties

- Purpose** Verify that the G.SHDSL interface properties are configured properly.
- Action** From operational mode, enter the **show interfaces at-1/0/0 extensive** command.

Release History Table	Release	Description
	15.1X49-D10	Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

- Related Documentation**
 - [Understanding Interfaces on page 3](#)
 - [SHDSL Interface Overview on page 145](#)
 - [G.SHDSL Mini-PIM Overview on page 146](#)
 - [G.SHDSL Mini-PIM Configuration Overview on page 148](#)

Example: Configuring the G.SHDSL Interface in EFM Mode

This example shows how to configure the G.SHDSL interface in Ethernet in the First Mile (EFM) mode on an SRX210 device, but it applies to the SRX220, SRX240, and SRX550 devices as well.



NOTE: Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

- [Requirements on page 170](#)
- [Overview and Topology on page 170](#)
- [Configuration on page 171](#)
- [Verification on page 175](#)

Requirements

This example uses the following hardware and software components:

- An SRX210 device
- Junos OS Release 12.1X44-D10 or later

Before you begin:

- Configure the network interfaces as necessary. See [“Understanding Ethernet Interfaces” on page 259](#).
- Install the G.SHDSL Mini-PIM in the first slot of the SRX210 chassis.
- Connect the SRX210 device to an EFM supported IP DSLAM.

Overview and Topology

In this example, 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.

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.

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 14 on page 171 shows the topology for the G.SHDSL Mini-PIM operating in EFM mode.

Figure 14: G.SHDSL Mini-PIM Operating in EFM Mode




Table 30 on page 171 lists the operating wire mode for EFM and its corresponding CLI code.

Table 30: Operating Wire Mode for EFM

Wire Mode Configuration	CLI Code
EFM Configuration	set chassis fpc 1 pic 0 shdsl pic-mode efm



NOTE: When PIC mode is set to EFM, an interface called pt-1/0/0 is created.

Configuration

- [Configuring a Basic G.SHDSL Interface in EFM PIC Mode on page 171](#)
- [Configuring PPPoE and VLAN for the G.SHDSL EFM Interface on page 173](#)

Configuring a Basic G.SHDSL Interface in EFM PIC Mode

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 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
```

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 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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** and **show chassis fpc 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 pt-1/0/0
shdsl-options {
  annex annex-g;
  line-rate 5696;
}
unit 0 {
  family inet {
    address 10.10.10.1/24;
  }
}
[edit]
user@host# show chassis fpc 1
pic 0 {
  shdsl {
    pic-mode efm;
  }
}
```

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

Configuring PPPoE and VLAN for the G.SHDSL EFM 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.



NOTE: In this configuration, we use PAP as the authentication mechanism. If Broadband Remote Access Server (BRAS) uses CHAP, PAP configuration should be replaced with CHAP.

```
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
```

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 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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** and **show interfaces pp0** 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 pt-1/0/0
vlan-tagging;
unit 0 {
    encapsulation ppp-over-ether;
    vlan-id 99;
}
[edit]
user@host# show interfaces pp0
unit 0 {
    ppp-options {
        pap {
            access-profile pap_prof;
            local-name srx-210;
            local-password "$9$0tLw1SeN-woJDSr-wY2GU69Cp1RSre";
            passive;
        }
    }
    pppoe-options {
        underlying-interface pt-1/0/0.0;
        auto-reconnect 120;
        client;
    }
    family inet {
        negotiate-address;
    }
}
```

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

Verification

Verifying G.SHDSL Interface Properties

Purpose Verify that the G.SHDSL interface properties are configured properly.

Action From operational mode, enter the **show interfaces pt-1/0/0 extensive** command.

```
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)
    Framing mode     : EFM
    PAF Status       : Active
    Dying gasp       : Enabled
    Framing 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
      0 best-effort           95          21644800    95          0          low
none
      3 network-control       5           1139200     5           0          low
none

```

Meaning 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 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 page (*Interfaces*>*interface-name*).
- 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 expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics *interface-name*** command.
- The following information is displayed for each line pair:

No SHDSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm.

 - **LOSW**—Loss of sync word. A message ID was sent.
 - **LOSD**—Loss of signal was detected at the remote application interface.
 - **ES**—Errored seconds. One or more cyclic redundancy check (CRC) anomalies were detected.
 - **SES**—Severely errored seconds. At least 50 CRC anomalies were detected.
 - **UAS**—Unavailable seconds. An interval has occurred during which one or more LOSW defects were detected.

Examine the SHDSL interface status:

- **Line termination**—SHDSL transceiver unit—remote (STU—R). (Only customer premises equipment is supported.)
- **Annex**—Either Annex A or Annex B. Annex A is supported in North America, and Annex B is supported in Europe.
- **Line mode**—SHDSL mode configured on the G.SHDSL interface pair, and it should be two-wire.
- **Modem status**—Data. Sending or receiving data.
- **Bit rate (kbps)**—Data transfer speed on the SHDSL interface.
- **Last fail code**—Code for the last interface failure.
- **Framer mode**—ATM framer mode of the underlying interface.
- **PAF Status**—Either Active/Inactive depending upon whether link added to EFM group or not.
-

Examine the operational statistics for a SHDSL interface.

- **Loop attenuation (dB)**—Reduction in signal strength.
- **Transmit power (dB)**—Amount of SHDSL.
- **SNR sampling (dB)**—Signal-to-noise ratio at a receiver point.
- **CRC errors**—Number of cyclic redundancy check errors.
- **SEGA errors**—Number of segment anomaly errors. A regenerator operating on a segment received corrupted data.
- **LOSW errors**—Number of loss of signal defect errors. Three or more consecutively received frames contained one or more errors in the framing bits.

Release History Table

Release	Description
15.1X49-D10	Starting in Junos OS Release 15.1X49-D10 SHDSL interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Related Documentation

- [SHDSL Interface Overview on page 145](#)
- [G.SHDSL Mini-PIM Overview on page 146](#)
- [G.SHDSL Mini-PIM Configuration Overview on page 148](#)
- [Example: Configuring the G.SHDSL Interface on page 150](#)
- [Example: Configuring the G.SHDSL Interface on SRX Series Devices on page 158](#)

CHAPTER 12

Configuring VDSL2 Interfaces

- [VDSL2 Interface Technology Overview on page 181](#)
- [VDSL2 Network Deployment Topology on page 182](#)
- [VDSL2 Interface Support on SRX Series Devices on page 184](#)
- [Example: Configuring VDSL2 Interfaces in ADSL Mode \(Basic\) on page 187](#)
- [Example: Configuring VDSL2 Interfaces in ADSL Mode \(Detail\) on page 193](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 220](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 227](#)
- [Upgrading the VDSL PIC Firmware on page 253](#)

VDSL2 Interface Technology Overview

Very-high-bit-rate digital subscriber line (VDSL) technology is part of the xDSL family of modem technologies that provide faster data transmission over a single flat untwisted or twisted pair of copper wires. The VDSL lines connect service provider networks and customer sites to provide high bandwidth applications (triple-play services) such as high-speed Internet access, telephone services like VoIP, high-definition TV (HDTV), and interactive gaming services over a single connection.

VDSL2 is an enhancement to G.993.1 (VDSL) and permits the transmission of asymmetric (half-duplex) and symmetric (full-duplex) aggregate data rates up to 100 Mbps on short copper loops using a bandwidth up to 30 MHz. The VDSL2 technology is based on the ITU-T G.993.2 (VDSL2) standard, which is the International Telecommunication Union standard describing a data transmission method for VDSL2 transceivers.

The VDSL2 uses discrete multitone (DMT) modulation. DMT is a method of separating a digital subscriber line signal so that the usable frequency range is separated into 256 frequency bands (or channels) of 4.3125 KHz each. The DMT uses the Fast Fourier Transform (FFT) algorithm for demodulation or modulation for increased speed.

VDSL2 interface supports Packet Transfer Mode (PTM). The PTM mode transports packets (IP, PPP, Ethernet, MPLS, and so on) over DSL links as an alternative to using Asynchronous Transfer Mode (ATM). PTM is based on the Ethernet in the First Mile (EFM) IEEE802.3ah standard.

VDSL2 provides backward compatibility with ADSL, ADSL2, and ADSL2+ because this technology is based on both the VDSL1-DMT and ADSL2/ADSL2+ recommendations.

VDSL2 Vectoring Overview

Starting in Junos OS Release 15.1X49-D50, VDSL2 vectoring is supported. Vectoring is a transmission method that employs the coordination of line signals that reduce crosstalk levels and improve performance. It is based on the concept of noise cancellation, like noise-cancelling headphones. The ITU-T G.993.5 standard, "Self-FEXT Cancellation (Vectoring) for Use with VDSL2 Transceivers," also known as G.vector, describes vectoring for VDSL2.

The scope of Recommendation ITU-T G.993.5 is specifically limited to the self-FEXT (far-end crosstalk) cancellation in the downstream and upstream directions. The FEXT generated by a group of near-end transceivers and interfering with the far-end transceivers of that same group is canceled. This cancellation takes place between VDSL2 transceivers, not necessarily of the same profile.

Release History Table

Release	Description
15.1X49-D50	Starting in Junos OS Release 15.1X49-D50, VDSL2 vectoring is supported.

Related Documentation

- [VDSL2 Network Deployment Topology on page 182](#)
- [VDSL2 Interface Support on SRX Series Devices on page 184](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 220](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 227](#)

VDSL2 Network Deployment Topology

In standard telephone cables of copper wires, voice signals use only a fraction of the available bandwidth. Like any other DSL technology, the VDSL2 technology utilizes the remaining capacity to carry the data and multimedia on the wire without interrupting the line's ability to carry voice signals.

This example depicts the typical VDSL2 network topology deployed using SRX Series Services Gateways.

A VDSL2 link between network devices is set up as follows:

1. Connect an end-user device such as a LAN, hub, or PC through an Ethernet interface to the customer premises equipment (CPE) (for example, an SRX Series device).
2. Connect the CPE to a DSLAM.

3. The VDSL2 interface uses either Gigabit Ethernet or fiber as second mile to connect to the Broadband Remote Access Server (B-RAS) as shown in [Figure 15 on page 183](#).
4. The ADSL interface uses either Gigabit Ethernet (in case of IP DSLAM) as the “second mile” to connect to the B-RAS or OC3/DS3 ATM as the second mile to connect the B-RAS as shown in [Figure 16 on page 183](#).

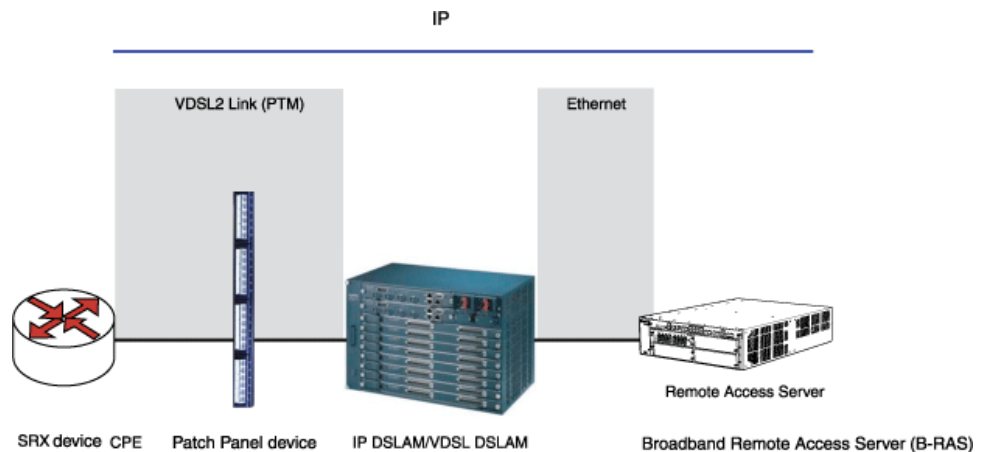


NOTE: The VDSL2 technology is backward compatible with ADSL. VDSL2 provides an ADSL interface in an ATM DSLAM topology and provides a VDSL2 interface in an IP or VDSL DSLAM topology.

The DSLAM accepts connections from many customers and aggregates them to a single, high-capacity connection to the Internet.


[Figure 15 on page 183](#) shows a typical VDSL2 network topology.

Figure 15: Typical VDSL2 End-to-End Connectivity and Topology Diagram



[Figure 16 on page 183](#) shows a backward-compatible ADSL topology using ATM DSLAM.

Figure 16: Backward-Compatible ADSL Topology (ATM DSLAM)



- Related Documentation**
- [VDSL2 Interface Technology Overview on page 181](#)
 - [VDSL2 Interface Support on SRX Series Devices on page 184](#)
 - [Example: Configuring VDSL2 Interfaces \(Basic\) on page 220](#)
 - [Example: Configuring VDSL2 Interfaces \(Detail\) on page 227](#)

VDSL2 Interface Support on SRX Series Devices

The VDSL2 interface is supported on the SRX Series devices listed in [Table 31 on page 184](#). (Platform support depends on the Junos OS release in your installation.)

Table 31: VDSL2 Annex A and Annex B Features

Features	POTS	ISDN
Devices	Integrated VDSL Module (SRX110-POTS, SRX320-POTS) VDSL Mini-PIM (SRX210, SRX220, SRX240, SRX320, SRX340)	Integrated VDSL Module (SRX110-ISDN, SRX320-ISDN)
Supported annex operating modes	Annex A and Annex B*	Annex B
Supported Bandplans	997/998	998
Supported standards	ITU-T G.993.2 and ITU-T G.993.5 (VDSL2)	ITU-T G.993.2 and ITU-T G.993.5 (VDSL2)
Used in	North American network implementations	European network implementations
ADSL backward compatibility	ADSL G992.5-A (ADSL Annex A)	ADSL G992.5-B (ADSL Annex B)

* Annex B support is not available on VDSL2 Mini-PIMs.

VDSL2 Interface Compatibility with ADSL Interfaces

VDSL2 interfaces on SRX Series devices are backward compatible with most ADSL interface standards. The VDSL2 interface uses Ethernet in the First Mile (EFM) mode or Packet Transfer Mode (PTM) and uses the named interface `pt-1/0/0`. In ADSL fallback mode, VDSL2 operates on the ATM encapsulation interface in the first mile and uses the named interface `at-1/0/0`.

**NOTE:**

- The VDSL2 interface has backward compatibility with ADSL/ADSL2/ADSL2+. The VDSL2 interface is represented by the `pt` interface when configured to function as VDSL2, and the ADSL interface is represented by the `at` interface when configured to function as ADSL.
- On VDSL2 interfaces, by default the `pt-1/0/0` interface is created when there is no configuration already created for either the `pt-1/0/0` or the `at-1/0/0` interface.



NOTE: It requires around 60 seconds to switch from VDSL2 to ADSL or from ADSL to VDSL2 operating modes.

Table 32 on page 185 lists VDSL2 operating modes and their backward compatibility with ADSL interface standards.

Table 32: VDSL2 Operating Mode Backward Compatibility with ADSL

VDSL2 Annex Type	Operating Modes	Description
VDSL2 Annex A interface (ADSL modes for Annex A only)	auto	Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex A, the ADSL interface uses either ANSI T1.413 Issue II mode or ITU G.992.1 mode. NOTE: Automatic (auto) operating mode does not work when the DSLAM located at the central office is operating at ADSL2+ Annex M mode.
	ansi-dmt	Configures the ADSL interface to use ANSI T1.413 Issue II mode.
	itu-dmt	Configures the ADSL interface to use ITU G.992.1 mode.
	itu-dmt-bis	Configures the ADSL interface to use ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM.
	adsl2plus	Configures the ADSL interface to use ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM.

Table 32: VDSL2 Operating Mode Backward Compatibility with ADSL (continued)

VDSL2 Annex Type	Operating Modes	Description
VDSL2 Annex B interface (ADSL modes for Annex B only)	auto	Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex B, the ADSL interface trains in ITU G.992.1 mode.
	itu-dmt	Configures the ADSL interface to use ITU G.992.1 mode.
	itu-dmt-bis	Configures the ADSL interface to use ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM.
	adsl2plus	Configures the ADSL interface to use ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM.
	itu-annexb-ur2	Configures the ADSL line to use G.992.1 Deutsche Telekom UR-2 mode.



NOTE: On SRX210, SRX220, and SRX240 devices, every time the VDSL2 Mini-PIM is restarted in the ADSL mode, the first packet passing through the Mini-PIM is dropped.

VDSL2 Interfaces Supported Profiles

A profile is a table that contains a list of preconfigured VDSL2 settings.

[Table 33 on page 186](#) lists the different profiles supported on the VDSL2 interfaces and their properties.

Table 33: Supported Profiles on the VDSL2 Interfaces

Profiles	Data Rate
8a	50
8b	50
8c	50
8d	50
12a	68
12b	68
17a	100

Table 33: Supported Profiles on the VDSL2 Interfaces (continued)

Profiles	Data Rate
Auto	Negotiated (based on operating mode)

VDSL2 Interfaces Supported Features

The following features are supported on the VDSL2 interfaces:

- ADSL/ADSL2/ADSL2+ backward compatibility with Annex A, Annex M support
- PTM or EFM (802.3ah) support
- Operation, Administration, and Maintenance (OAM) support for ADSL/ADSL2/ADSL2+ mode
- ATM quality of service (QoS) (supported only when the VDSL2 Mini-PIM is operating in ADSL2 mode)
- Multilink Point-to-Point Protocol (MLPPP) (supported only when the VDSL2 Mini-PIM is operating in ADSL2 mode)
- MTU size of 1514 bytes (maximum) in VDSL2 mode and 1496 bytes in ADSL mode.
- Support for maximum of 10 permanent virtual connections (PVCs) (only in ADSL/ADSL2/ADSL2+ mode)
- Dying gasp support (ADSL and VDSL2 mode)



NOTE: On SRX210 or SRX320 devices with VDSL2, ATM CoS VBR-related functionality cannot be tested.

Related Documentation

- [VDSL2 Interface Technology Overview on page 181](#)
- [VDSL2 Network Deployment Topology on page 182](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 220](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 227](#)

Example: Configuring VDSL2 Interfaces in ADSL Mode (Basic)

This example shows how to configure the integrated VDSL2 interfaces for SRX320 (Annex B) in ADSL backward compatible mode.

- [Requirements on page 188](#)
- [Overview on page 188](#)
- [Configuration on page 188](#)
- [Verifying the Configuration on page 189](#)

Requirements

Before you begin:

- Set up and perform initial configuration on the SRX Series devices.
- Connect the SRX320 device to a DSLAM
- Establish basic connectivity. See the *Quick Start Guide* for your device for factory default settings.
- On VDSL2 interfaces, by default the `pt-1/0/0` interface is created when there is no configuration already created for either the `pt-1/0/0` or the `at-1/0/0` interface. You can switch to ADSL mode by just configuring `at-1/0/0`. If the configurations are already created for `pt-1/0/0` or `at-1/0/0`, then you need to 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`.

Overview

In this example, you create a VDSL2 interface called `pt-1/0/0`, specify the type of encapsulation, and set the VDSL2 profile to auto.

Configuration

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 fe-0/0/3 unit 0 family inet address 10.10.10.1/24
set interfaces at-1/0/0 atm-options vpi 0
set interfaces at-1/0/0 dsl-options operating-mode auto
set interfaces at-1/0/0 unit 0 vci 0.33
```

Step-by-Step Procedure

To configure the VDSL2 interfaces for the SRX320 in ADSL backward compatible mode:

1. Set operating mode.

```
[edit]
user@host# user@host# set interfaces at-1/0/0 dsl-options operating-mode auto
```
2. Configure the ATM VPI option

```
[edit]
user@host# set interfaces at-1/0/0 atm-options vpi 0
```
3. Set the ATM VCI option.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 vci 0.33
```

4. Configure the IP address for the interface.

[edit]

user@host# set interfaces fe-0/0/3 unit 0 family inet address 10.10.10.1/24

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

Verifying the Configuration

Confirm that the configuration is working properly.

Verifying the Configuration

Purpose Verify the command output.

Action From operational mode, enter the **show interfaces at-1/0/0 extensive** command.

```
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 513, Generation: 175
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode,
  Speed: ADSL2+
  Speed: 1573kbps, 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: 00:1f:12:e4:df:20
  Last flapped   : 2011-05-25 05:58:32 PDT (00:02:54 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 channel errors: 0, L2 mismatch timeouts: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, 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      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
ADSL alarms       : None
ADSL defects      : None
ADSL media:
Seconds          Count  State
LOF              55      0 OK
LOS              55      0 OK
LOM              0       0 OK
LOP              0       0 OK
LOCDI            0       0 OK
LOCDNI           55      0 OK
ADSL status:
Modem status      : Showtime (Adsl2plus)
DSL mode          : Auto Annex B   Last fail code: None
Subfunction       : 0x00
Seconds in showtime: 173
ADSL Chipset Information:
Vendor Country    : ATU-R          ATU-C
Vendor ID         : 0xb5           0xb5
Vendor Specific   : BDCM           BDCM
Vendor Specific   : 0x9385         0x9395
ADSL Statistics:
Attenuation (dB)  : ATU-R          ATU-C
Capacity used (%) : 1.5            0.0
Noise margin (dB) : 0              0
Output power (dBm) : 8.5           9.0
                  : 6.5           9.0

Bit rate (kbps)   : Interleave    Fast Interleave    Fast
CRC               : 24681         0      1573         0
FEC               : 0             0       0            0
HEC               : 0             0       0            0
Received cells    : 278817900     0
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: 0, Output cell count: 0,
Output idle cell count: 0, Output VC queue drops: 0, Input 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
CoS information:
Direction : Output
CoS transmit queue
Limit      Bandwidth      Buffer Priority
%          bps           %          usec
0 best-effort 95      1494350 95      0      low
none
3 network-control 5      78650 5      0      low
none

```

```

Logical interface at-1/0/0.0 (Index 73) (SNMP ifIndex 533) (Generation 157)
Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-SNAP
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: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmp igmp ldp msdp nhrp
ospf pgm pim rip 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
VCI 0.33
Flags: Active
Total down time: 0 sec, Last down: Never
ATM per-VC transmit statistics:
  Tail queue packet drops: 0
Traffic statistics:
  Input bytes : 0
  Output bytes : 0

```

```

        Input  packets:           0
        Output packets:           0

Logical interface at-1/0/0.32767 (Index 74) (SNMP ifIndex 534)
(Generation 158)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0
  Encapsulation: ATM-VCMUX
  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
  Security: Zone: HOST
  Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
  ospf pgm pim rip router-discovery rsvp sap vrrp
  Flow Statistics :
  Flow Input statistics :
    Self packets :           0
    ICMP packets :           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
  VCI 0.4
  Flags: Active
  Total down time: 0 sec, Last down: Never
  ATM per-VC transmit statistics:
  Tail queue packet drops: 0
  Traffic statistics:
    Input  bytes  :           0
    Output bytes  :           0

```



```

Input  packets:          0
Output packets:          0

```

The output shows a summary of VDSL2 interface. Verify the following information:

- Status of interface at-1/0/0 is displayed as **Physical link is Up**.
- Modem status is displayed as **Showtime (Adsl2plus)**.
- Time in seconds during which the interface stayed up is displayed as **Seconds in showtime**.
- ADSL profile of the DSLAM is displayed as **Annex B**.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [VDSL2 Interface Technology Overview on page 181](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 227](#)
- [Example: Configuring VDSL2 Interfaces in ADSL Mode \(Detail\) on page 193](#)

Example: Configuring VDSL2 Interfaces in ADSL Mode (Detail)

This example shows how to configure ADSL Interfaces for SRX Series devices.

This example uses VDSL2 Mini-PIM installed on SRX320 devices. The information is also applicable to SRX340 devices (with VDSL2 Mini-PIMs).

- [Requirements on page 193](#)
- [Overview on page 193](#)
- [Configuration on page 194](#)
- [Verification on page 207](#)

Requirements

Before you begin:

- Install Junos OS Release 10.1 or later for the SRX Series devices.
- Set up and perform initial configuration on the SRX Series device. See *Quick Start Guide* of your device for factory default settings.
- Install the VDSL2 Mini-PIM on the SRX320 device chassis.
- Ensure that the SRX320 device is connected to a DSLAM that supports VDSL2-to-ADSL fallback.

Overview

In this example, you configure the ADSL interface for end-to-end data path. Then you configure PPPoA on the at-1/0/0 interface with a negotiated IP address and either PAP authentication or CHAP authentication. You also 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.

Finally, you configure PPPoE on the at-1/0/0 interface with a negotiated IP address and either PAP authentication or CHAP authentication.

Configuration

- [Configuring the ADSL Interface for End-to-End Data Path on page 194](#)
- [Configuring PPPoA on the at-1/0/0 Interface with Negotiated IP and PAP Authentication on page 195](#)
- [Configuring PPPoA on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication on page 197](#)
- [Configuring PPPoA on the at-1/0/0 Interface with Static IP and PAP Authentication on page 198](#)
- [Configuring PPPoA on the at-1/0/0 Interface with Static IP and CHAP Authentication on page 199](#)
- [Configuring PPPoA on the at-1/0/0 Interface with Unnumbered IP and PAP Authentication on page 201](#)
- [Configuring PPPoA on the at-1/0/0 Interface with Unnumbered IP and CHAP Authentication on page 202](#)
- [Configuring PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and PAP Authentication on page 204](#)
- [Configuring PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication on page 206](#)

Configuring the ADSL Interface for End-to-End Data Path

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-1/0/0 encapsulation atm-pvc atm-options vpi 2
set interfaces at-1/0/0 dsl-options operating-mode itu-dmt
set interfaces at-1/0/0 unit 0 encapsulation atm-snap vci 2.119 family inet address
10.10.10.1/24
```

Step-by-Step Procedure

To configure the ADSL interface for end-to-end data path:

1. Delete any previous configurations.

[edit]
user@host# delete interfaces at-1/0/0
2. Specify the basic configuration for the ADSL interface.

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

```

user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 dsl-options operating-mode itu-dmt
user@host# set interfaces at-1/0/0 unit 0 encapsulation atm-snap
user@host# set interfaces at-1/0/0 unit 0 vci 2.119
user@host# set interfaces at-1/0/0 unit 0 family inet address 10.10.10.1/24

```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-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 at-1/0/0
encapsulation atm-pvc;
atm-options {
  vpi 2;
}
dsl-options {
  operating-mode itu-dmt;
}
encapsulation atm-snap;
vci 2.119;
family inet {
  address 10.10.10.1/24;
}
}

```

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

Configuring PPPoA on the at-1/0/0 Interface with Negotiated IP and PAP Authentication

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-1/0/0 encapsulation atm-pvc atm-options vpi 2
set interfaces at-1/0/0 dsl-options operating-mode auto
set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc vci 2.119
set interfaces at-1/0/0 unit 0 ppp-options pap access-profile jnpr local-name locky
  local-password india
set interfaces at-1/0/0 unit 0 family inet negotiate-address
set access profile jnpr client sringeri pap-password india

```

Step-by-Step Procedure To configure PPPoA on the at-1/0/0 interface with negotiated IP and PAP authentication:

1. Configure encapsulation and ATM options.

```

[edit]
user@host# set interfaces at-1/0/0 encapsulation atm-pvc
user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 dsl-options operating-mode auto
user@host# set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc

```

```
user@host# set interfaces at-1/0/0 unit 0 vci 2.119
```

2. Specify PPP options.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap access-profile jnpr
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap local-name locky
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap local-password india
```

3. Configure the negotiated IP address.

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

4. Configure the access profile.

```
[edit]
user@host# set access profile jnpr client sringeri pap-password india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show access profile jnpr** 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 at-1/0/0
encapsulation atm-pvc;
atm-options {
  vpi 2;
}
dsl-options {
  operating-mode auto;
}
unit 0 {
  encapsulation atm-ppp-llc;
  vci 2.119;
  ppp-options {
    pap {
      access-profile jnpr;
      local-name locky;
      local-password "$9$tm/auBEx7V2gJevWx"; ## SECRET-DATA
    }
  }
  family inet {
    negotiate-address;
  }
}
[edit]
user@host# show access profile jnpr
client sringeri pap-password "$9$FoPYn9peK8N-wRhSe"; ## SECRET-DATA
```

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

Configuring PPPoA on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication

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-1/0/0 encapsulation atm-pvc atm-options vpi 2
set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc vci 2.119
set interfaces at-1/0/0 unit 0 ppp-options chap access-profile jnpr local-name locky
set interfaces at-1/0/0 unit 0 family inet negotiate-address
set access profile jnpr client sringeri chap-secret india
```

Step-by-Step Procedure To configure PPPoA on the at-1/0/0 interface with negotiated IP and CHAP Authentication:

1. Configure encapsulation and ATM options.

```
[edit]
user@host# set interfaces at-1/0/0 encapsulation atm-pvc
user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc
user@host# set interfaces at-1/0/0 unit 0 vci 2.119
```

2. Specify PPP options.

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

3. Configure the negotiated IP address.

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

4. Configure the access profile.

```
[edit]
user@host# set access profile jnpr client sringeri chap-secret india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show access profile jnpr** 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 at-1/0/0
encapsulation atm-pvc;
atm-options {
  vpi 2;
```

```

}
unit 0 {
    encapsulation atm-ppp-llc;
    vci 2.119;
    ppp-options {
    chap {
        access-profile jnpr;
        local-name locky;
    }
    }
    family inet {
        negotiate-address;
    }
    }
[edit]
user@host# show access profile jnpr
client sringeri chap-secret "$9$qm5FIRSKvLAp0l"; ## SECRET-DATA

```

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

Configuring PPPoA on the at-1/0/0 Interface with Static IP and PAP Authentication

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-1/0/0 encapsulation atm-pvc atm-options vpi 2
set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc vci 2.119
set interfaces at-1/0/0 unit 0 ppp-options pap access-profile jnpr local-name locky
local-password india
set interfaces at-1/0/0 unit 0 family inet address 100.100.100.1/24
set access profile jnpr client sringeri pap-password india

```

Step-by-Step Procedure

To configure PPPoA on the at-1/0/0 interface with static IP and PAP authentication:

1. Configure encapsulation and ATM options.

```

[edit]
user@host# set interfaces at-1/0/0 encapsulation atm-pvc
user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc
user@host# set interfaces at-1/0/0 unit 0 vci 2.119

```

2. Specify PPP options.

```

[edit]
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap access-profile jnpr
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap local-name locky
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap local-password india

```

3. Configure the negotiated IP address.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet address 100.100.100.1/24
```

4. Configure the access profile.

```
[edit]
user@host# set access profile jnpr client sringeri pap-password india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show access profile jnpr** 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 at-1/0/0
encapsulation atm-pvc;
atm-options {
  vpi 2;
}
unit 0 {
  encapsulation atm-ppp-llc;
  vci 2.119;
  ppp-options {
  pap {
    access-profile jnpr;
    local-name locky;
    local-password "$9$GoDHmtpBhclFn/t"; ## SECRET-DATA
  }
  family inet {
    address 100.100.100.1/24;
  }
}
[edit]
user@host# show access profile jnpr
client sringeri pap-password "$9$p87c01h7Nbg4ZKM87"; ## SECRET-DATA
```

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

Configuring PPPoA on the at-1/0/0 Interface with Static IP and CHAP Authentication

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-1/0/0 encapsulation atm-pvc atm-options vpi 2
set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc vci 2.119
set interfaces at-1/0/0 unit 0 ppp-options chap access-profile jnpr local-name locky
set interfaces at-1/0/0 unit 0 family inet address 100.100.100.1/24
set access profile jnpr client sringeri chap-secret india
```

Step-by-Step Procedure To configure PPPoA on the at-1/0/0 interface with static IP and CHAP authentication:

1. Configure encapsulation and ATM options.

```
[edit]
user@host# set interfaces at-1/0/0 encapsulation atm-pvc
user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc
user@host# set interfaces at-1/0/0 unit 0 vci 2.119
```

2. Specify PPP options.

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

3. Configure the negotiated IP address.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet address 100.100.100.1/24
```

4. Configure the access profile.

```
[edit]
user@host# set access profile jnpr client sringeri chap-secret india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show access profile jnpr** 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 at-1/0/0
encapsulation atm-pvc;
atm-options {
  vpi 2;
}
unit 0 {
  encapsulation atm-ppp-llc;
  vci 2.119;
  ppp-options {
    chap {
      access-profile jnpr;
      local-name locky;
    }
  }
  family inet {
    address 100.100.100.1/24;
  }
}
[edit]
user@host# show access profile jnpr
client sringeri chap-secret "$9$mfQnEhrMWxp0BE"; ## SECRET-DATA
```


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

Configuring PPPoA on the at-1/0/0 Interface with Unnumbered IP and PAP Authentication

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-1/0/0 encapsulation atm-pvc atm-options vpi 2
set interfaces at-1/0/0 dsl-options operating-mode auto
set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc vci 2.119
set interfaces at-1/0/0 unit 0 ppp-options pap access-profile jnpr local-name locky
  local-password india
set interfaces at-1/0/0 unit 0 family inet unnumbered-address lo0.0 destination
  100.100.100.6
set interfaces lo0 unit 0 family inet address 100.100.100.20/32
set access profile jnpr client sringeri pap-password india
```

Step-by-Step Procedure

To configure PPPoA on the at-1/0/0 interface with unnumbered IP and PAP authentication:

1. Configure encapsulation and ATM options.

```
[edit]
user@host# set interfaces at-1/0/0 encapsulation atm-pvc
user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 dsl-options operating-mode auto
user@host# set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc
user@host# set interfaces at-1/0/0 unit 0 vci 2.119
```

2. Specify PPP options.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap access-profile jnpr
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap local-name locky
user@host# set interfaces at-1/0/0 unit 0 ppp-options pap local-password india
```

3. Configure the IP address, unnumbered IP address, and destination IP address.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces at-1/0/0 unit 0 family inet unnumbered-address
  destination 100.100.100.6
user@host# set interfaces lo0 unit 0 family inet address 100.100.100.20/32
```

4. Configure the access profile.

```
[edit]
user@host# set access profile jnpr client sringeri pap-password india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0**, **show interfaces lo0**, and **show access profile jnpr** 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 at-1/0/0
encapsulation atm-pvc;
  atm-options {
    vpi 2;
  }
  dsl-options {
    operating-mode auto;
  }
unit 0 {
  encapsulation atm-ppp-llc;
  vci 2.119;
  ppp-options {
  pap {
    access-profile jnpr;
    local-name locky;
    local-password "$9$LA7x-wHkPzF/aZUH"; ## SECRET-DATA
  }
  }
  family inet {
    unnumbered-address lo0.0 destination 100.100.100.6;
  }
}
[edit]
user@host# show interfaces lo0
unit 0 {
  family inet {
    address 100.100.100.20/32;
  }
}
[edit]
user@host# show access profile jnpr
client sringeri pap-password "$9$mSRclbwgZGiLxNb"; ## SECRET-DATA
```

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

Configuring PPPoA on the at-1/0/0 Interface with Unnumbered IP and CHAP Authentication

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-1/0/0 encapsulation atm-pvc atm-options vpi 2
set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc vci 2.119
set interfaces at-1/0/0 unit 0 ppp-options chap access-profile jnpr local-name locky
set interfaces at-1/0/0 unit 0 family inet unnumbered-address lo0.0 destination
100.100.100.6
```

```
set interfaces lo0 unit 0 family inet address 100.100.100.10/32
set access profile jnpr client sringeri chap-secret india
```

Step-by-Step Procedure To configure PPPoA on the at-1/0/0 interface with unnumbered IP and CHAP authentication:

1. Configure encapsulation and ATM-options.

```
[edit]
user@host# set interfaces at-1/0/0 encapsulation atm-pvc
user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 unit 0 encapsulation atm-ppp-llc
user@host# set interfaces at-1/0/0 unit 0 vci 2.119
```

2. Specify the PPP-options.

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

3. Configure the IP address, unnumbered IP address, and destination IP address.

```
[edit]
user@host# set interfaces at-1/0/0 unit 0 family inet unnumbered-address lo0.0
user@host# set interfaces at-1/0/0 unit 0 family inet unnumbered-address
destination 100.100.100.6
user@host# set interfaces lo0 unit 0 family inet address 100.100.100.10/32
```

4. Configure the access profile.

```
[edit]
user@host# set access profile jnpr client sringeri chap-secret india
```

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0**, **show interfaces lo0**, and **show access profile jnpr** 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 at-1/0/0
show interfaces at-1/0/0
  atm-options {
    vpi 2;
  }
  unit 0 {
    encapsulation atm-ppp-llc;
    vci 2.119;
    ppp-options {
      chap {
        access-profile jnpr;
        local-name locky;
      }
    }
  }
}
```

```

family inet {
    unnumbered-address lo0.0 destination 100.100.100.6;
}
[edit]
user@host# show interfaces lo0
unit 0 {
    family inet {
        address 100.100.100.10/32;
    }
}
[edit]
user@host# show access profile jnpr
client sringeri chap-secret "$9$.PT3REyvMXtuOR"; ## SECRET-DATA

```

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

Configuring PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and PAP Authentication

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-1/0/0 encapsulation ethernet-over-atm atm-options vpi 2
set interfaces at-1/0/0 unit 0 vci 2.119 encapsulation ppp-over-ether-over-atm-llc
set interfaces pp0 unit 0 ppp-options pap access-profile my_prf local-name purple
  local-password <password> passive
set interfaces pp0 unit 0 pppoe-options underlying-interface at-1/0/0.0 auto-reconnect
  120 client
set interfaces pp0 unit 0 family inet negotiate-address
set access profile my_prf authentication-order password
set access profile my_prf

```

Step-by-Step Procedure

To configure PPPoE over ATM on the at-1/0/0 interface with negotiated IP and PAP authentication:

1. Configure encapsulation and ATM options.

```

[edit]
user@host# set interfaces at-1/0/0 encapsulation ethernet-over-atm
user@host# set interfaces at-1/0/0 atm-options vpi 2
user@host# set interfaces at-1/0/0 unit 0 vci 2.119
user@host# set interfaces at-1/0/0 unit 0 encapsulation
  ppp-over-ether-over-atm-llc

```

2. Specify PPP options.

```

[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. Specify PPPoE options.

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

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

5. Configure the access profile.

```
[edit]
user@host# set access profile my_prf authentication-order password
user@host# set access profile my_prf
```

Results From configuration mode, confirm your configuration by entering the **set access profile my_prf**, **show access profile my_prf**, and **show interfaces pp0** 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 at-1/0/0
encapsulation ethernet-over-atm;
atm-options {
    vpi 2;
}
unit 0 {
    encapsulation ppp-over-ether-over-atm-llc;
    vci 2.119;
}
[edit]
user@host# show access profile my_prf
authentication-order password;
[edit]
user@host# show interfaces pp0
unit 0 {
    ppp-options {
    pap {
        access-profile my_prf;
        local-name purple;
        local-password "$9$YkgoZTQn9CuZU69A0hcdbsYoGikP"; ## SECRET-DATA
        passive;
    }
    }
    pppoe-options {
        underlying-interface at-1/0/0.0;
        auto-reconnect 120;
        client;
    }
    family inet {
        negotiate-address;
```

```
}  
}
```

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

Configuring PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication

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-1/0/0 encapsulation ethernet-over-atm atm-options vpi 2  
set interfaces at-1/0/0 unit 0 vci 2.119 encapsulation ppp-over-ether-over-atm-llc  
set interfaces pp0 unit 0 ppp-options chap default-chap-secret <password> local-name  
  purple passive  
set interfaces pp0 unit 0 pppoe-options underlying-interface at-1/0/0.0 auto-reconnect  
  120 client  
set interfaces pp0 unit 0 family inet negotiate-address
```

Step-by-Step Procedure

To configure PPPoE over ATM on the at-1/0/0 interface with negotiated IP and CHAP authentication:

1. Configure encapsulation and ATM options.

```
[edit]  
user@host# set interfaces at-1/0/0 encapsulation ethernet-over-atm  
user@host# set interfaces at-1/0/0 atm-options vpi 2  
user@host# set interfaces at-1/0/0 unit 0 vci 2.119  
user@host# set interfaces at-1/0/0 unit 0 encapsulation  
  ppp-over-ether-over-atm-llc
```

2. Specify PPP options.

```
[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. Specify PPPoE options.

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

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

Results From configuration mode, confirm your configuration by entering the **show interfaces at-1/0/0** and **show interfaces pp0** 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 at-1/0/0
encapsulation ethernet-over-atm;
atm-options {
  vpi 2;
}
unit 0 {
  encapsulation ppp-over-ether-over-atm-llc;
  vci 2.119;
}
[edit]
user@host# show interfaces pp0
unit 0 {
  ppp-options {
    chap {
      default-chap-secret "$9$QQCIFn9cSeMWx9AKM87sYmfTQnCuOR"; ##
      SECRET-D ATA
      local-name purple;
      passive;
    }
  }
  pppoe-options {
    underlying-interface at-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    negotiate-address;
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the ADSL Interface for End-to-End Data Path on page 208](#)
- [Verifying PPPoA on the at-1/0/0 Interface with Negotiated IP and PAP Authentication on page 209](#)
- [Verifying PPPoA on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication on page 210](#)
- [Verifying PPPoA on the at-1/0/0 Interface with Static IP and PAP Authentication on page 212](#)
- [Verifying PPPoA on the at-1/0/0 Interface with Static IP and CHAP Authentication on page 213](#)
- [Verifying PPPoA on the at-1/0/0 Interface with Unnumbered IP and PAP Authentication on page 215](#)

- [Verifying PPPoA on the at-1/0/0 Interface with Unnumbered IP and CHAP Authentication on page 216](#)
- [Verifying PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and PAP Authentication on page 218](#)
- [Verifying PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication on page 219](#)

Verifying the ADSL Interface for End-to-End Data Path

Purpose Verify the interface status and traffic statistics.

Action From operational mode, enter the **show interface at-1/0/0 terse** and **show interfaces at-1/0/0** commands.

```
user@host> 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    10.10.10.1/24
  at-1/0/0.32767        up   up

[edit]
user@host# run ping 10.10.10.2 count 1000 rapid
PING 10.10.10.2 (10.10.10.2): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 10.10.10.2 ping statistics ---
1000 packets transmitted, 1000 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.141/9.356/58.347/3.940 ms

[edit]
user@host#

user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 504
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode,
  Speed: ADSL
  Speed: 832kbps, Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:b1:7e:85:84:ff
  Last flapped   : 2009-10-28 02:14:45 PDT (00:09:54 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  ADSL alarms    : None
  ADSL defects   : None
  ADSL status:
    Modem status  : Showtime (Itu-dmt)
    DSL mode      : Itu-dmt   Annex A
    Last fail code: None
    Subfunction   : 0x00
    Seconds in showtime : 596

  Logical interface at-1/0/0.0 (Index 69) (SNMP ifIndex 523)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-SNAP
```



```

Input packets : 1000
Output packets: 1000
Security: Zone: Null
Protocol inet, MTU: 1456
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
VCI 2.119
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 1000
  Output packets: 1000

Logical interface at-1/0/0.32767 (Index 70) (SNMP ifIndex 525)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0
  Encapsulation: ATM-VCMUX
  Input packets : 0
  Output packets: 0
  Security: Zone: Null
VCI 2.4
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 0
  Output packets: 0

```

Verifying PPPoA on the at-1/0/0 Interface with Negotiated IP and PAP Authentication

Purpose Verify the interface status and end-to-end data path connectivity.

Action From operational mode, enter the **show interfaces at-1/0/0** and **show interfaces at-1/0/0 terse** commands.

```

user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 504
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode, Speed: ADSL

Speed: 832kbps, Loopback: None
Device flags   : Present Running
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 00:b1:7e:85:84:ff
Last flapped   : 2009-10-28 02:39:14 PDT (00:09:29 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 80 bps (0 pps)
ADSL alarms    : None
ADSL defects   : None
ADSL status:
  Modem status  : Showtime (Itu-dmt)
  DSL mode      : Auto Annex A
  Last fail code: None
  Subfunction   : 0x00
  Seconds in showtime : 571

Logical interface at-1/0/0.0 (Index 69) (SNMP ifIndex 523)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-PPP-LLC

```

```

    Input packets : 2
    Output packets: 2
    Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Keepalive: Input: 8 (00:00:01 ago), Output: 9 (00:00:03 ago)
    LCP state: Opened
    NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
    CHAP state: Closed
    PAP state: Success
    Security: Zone: Null
    Protocol inet, MTU: 1486
    Flags: Negotiate-Address
    Addresses, Flags: Kernel Is-Preferred Is-Primary
    Destination: 100.100.100.6, Local: 100.100.100.1
    VCI 2.119
    Flags: Active
    Total down time: 0 sec, Last down: Never
    Input packets : 2
    Output packets: 2

Logical interface at-1/0/0.32767 (Index 70) (SNMP ifIndex 525)
Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0 Encapsulation: ATM-VCMUX

    Input packets : 0
    Output packets: 0
    Security: Zone: Null
    VCI 2.4
    Flags: Active
    Total down time: 0 sec, Last down: Never
    Input packets : 0
    Output packets: 0

user@host> 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    100.100.100.1        --> 100.100.100.6
at-1/0/0.32767  up    up

[edit]
user@host# run ping 100.100.100.6 count 100 rapid
PING 100.100.100.6 (100.100.100.6): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 100.100.100.6 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.056/8.501/14.194/1.787 ms

```

Verifying PPPoA on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication

Purpose Verify the interface output and end-to-end data path connectivity.

Action From operational mode, enter the **show interfaces at-1/0/0** and **show interfaces at-1/0/0 terse** commands.

```
user@host> show interfaces at-1/0/0
```

```

Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 504
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode, Speed: ADSL

  Speed: 832kbps, Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:b1:7e:85:84:ff
  Last flapped   : 2009-10-28 02:39:14 PDT (00:01:37 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 80 bps (0 pps)
  ADSL alarms    : None
  ADSL defects   : None
  ADSL status:
    Modem status : Showtime (Itu-dmt)
    DSL mode      : Auto Annex A
    Last fail code: None
    Subfunction   : 0x00
    Seconds in showtime : 97

Logical interface at-1/0/0.0 (Index 71) (SNMP ifIndex 523)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-PPP-LLC
  Input packets : 26
  Output packets: 29
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 10 (00:00:02 ago), Output: 8 (00:00:06 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: 1486
  Flags: Negotiate-Address
  Addresses, Flags: Kernel Is-Preferred Is-Primary
    Destination: 100.100.100.6, Local: 100.100.100.1
  VCI 2.119
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 26
  Output packets: 29

Logical interface at-1/0/0.32767 (Index 70) (SNMP ifIndex 525)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0 Encapsulation: ATM-VCMUX

  Input packets : 0
  Output packets: 0
  Security: Zone: Null
  VCI 2.4
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 0
  Output packets: 0

user@host> 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   100.100.100.1  --> 100.100.100.6
at-1/0/0.32767 up    up

```

```
[edit]
user@host# run ping 100.100.100.6 count 100 rapid
PING 100.100.100.6 (100.100.100.6): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 100.100.100.6 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.231/9.167/58.852/5.716 ms
```

Verifying PPPoA on the at-1/0/0 Interface with Static IP and PAP Authentication

Purpose Verify the interface status and end-to-end data path testing.

Action From operational mode, enter the **show interfaces at-1/0/0** and **show interfaces at-1/0/0 terse** commands.

```
user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 504
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode, Speed: ADSL

  Speed: 832kbps, Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:b1:7e:85:84:ff
  Last flapped   : 2009-10-28 22:18:50 PDT (00:10:26 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 80 bps (0 pps)
  ADSL alarms    : None
  ADSL defects   : None
  ADSL status:
    Modem status : Showtime (Itu-dmt)
    DSL mode      :      Auto      Annex A
    Last fail code: None
    Subfunction   : 0x00
    Seconds in showtime : 624

  Logical interface at-1/0/0.0 (Index 73) (SNMP ifIndex 523)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-PPP-LLC
    Input packets : 28
    Output packets: 29
    Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Keepalive: Input: 2 (00:00:01 ago), Output: 1 (00:00:09 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: 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: 1486
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 100.100.100/24, Local: 100.100.100.10, Broadcast:
```

```

100.100.100.255
  VCI 2.119
    Flags: Active
    Total down time: 0 sec, Last down: Never
    Input packets : 28
    Output packets: 29

Logical interface at-1/0/0.32767 (Index 72) (SNMP ifIndex 525)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0 Encapsulation: ATM-VCMUX

  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
  VCI 2.4
    Flags: Active
    Total down time: 0 sec, Last down: Never
    Input packets : 0
    Output packets: 0

user@host> 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    100.100.100.10/24
at-1/0/0.32767      up    up

[edit]
user@host# run ping 100.100.100.6 count 100 rapid
PING 100.100.100.6 (100.100.100.6): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 100.100.100.6 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.698/10.296/61.622/5.856 ms

```

Verifying PPPoA on the at-1/0/0 Interface with Static IP and CHAP Authentication

Purpose Verify the interface status and end-to-end data path testing.

Action From operational mode, enter the **show interfaces at-1/0/0** and **show interfaces at-1/0/0 terse** commands.

```

user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 504
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode, Speed: ADSL

  Speed: 832kbps, Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:b1:7e:85:84:ff
  Last flapped   : 2009-10-28 22:18:50 PDT (00:05:17 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  ADSL alarms    : None

```

```

ADSL defects : None
ADSL status:
  Modem status : Showtime (Itu-dmt)
  DSL mode      : Auto Annex A
  Last fail code: None
  Subfunction   : 0x00
  Seconds in showtime : 316

Logical interface at-1/0/0.0 (Index 71) (SNMP ifIndex 523)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-PPP-LLC
  Input packets : 46
  Output packets: 88
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 18 (00:00:04 ago), Output: 17 (00:00:08 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: 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: 1486
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 100.100.100/24, Local: 100.100.100.1, Broadcast:
100.100.100.255
  VCI 2.119
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 46
  Output packets: 88

Logical interface at-1/0/0.32767 (Index 72) (SNMP ifIndex 525)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0 Encapsulation: ATM-VCMUX

  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
  VCI 2.4
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 0
  Output packets: 0

user@host> 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    100.100.100.1/24
at-1/0/0.32767      up    up

[edit]
user@host# run ping 100.100.100.6 count 100 rapid
PING 100.100.100.6 (100.100.100.6): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 100.100.100.6 ping statistics ---

```

100 packets transmitted, 100 packets received, 0% packet loss
 round-trip min/avg/max/stddev = 7.787/9.300/15.081/2.023 ms

Verifying PPPoA on the at-1/0/0 Interface with Unnumbered IP and PAP Authentication

Purpose Verify the interface status and end-to-end data path testing.

Action From operational mode, enter the **show interfaces at-1/0/0** and **show interfaces at-1/0/0 terse** commands.

```
user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 504
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode, Speed: ADSL

  Speed: 832kbps, Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:b1:7e:85:84:ff
  Last flapped   : 2009-10-28 22:18:50 PDT (00:19:19 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  ADSL alarms    : None
  ADSL defects   : None
  ADSL status:
    Modem status : Showtime (Itu-dmt)
    DSL mode      : Auto Annex A
    Last fail code: None
    Subfunction   : 0x00
    Seconds in showtime : 1158

  Logical interface at-1/0/0.0 (Index 73) (SNMP ifIndex 523)
    Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-PPP-LLC
    Input packets : 441
    Output packets: 342
    Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Keepalive: Input: 53 (00:00:06 ago), Output: 55 (00:00:05 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: 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: 1486
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 100.100.100/24, Local: 100.100.100.20, Broadcast:
    100.100.100.255
    VCI 2.119
    Flags: Active
    Total down time: 0 sec, Last down: Never
    Input packets : 441
    Output packets: 342
```

```

Logical interface at-1/0/0.32767 (Index 72) (SNMP ifIndex 525)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0 Encapsulation: ATM-VCMUX

  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
  VCI 2.4
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 0
  Output packets: 0

user@host> show interfaces at-1/0/0 terse
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    100.100.100.20  --> 100.100.100.6
at-1/0/0.32767      up    up

[edit]
user@host# run ping 100.100.100.6 count 100 rapid
PING 100.100.100.6 (100.100.100.6): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 100.100.100.6 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.917/10.164/56.428/5.340 ms

```

Verifying PPPoA on the at-1/0/0 Interface with Unnumbered IP and CHAP Authentication

Purpose Verify the interface status and end-to-end data path connectivity.

Action From operational mode, enter the **show interfaces at-1/0/0** and **show interfaces at-1/0/0 terse** commands.

```

user@host> show interfaces at-1/0/0
Physical interface: at-1/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 504
  Link-level type: ATM-PVC, MTU: 1496, Clocking: Internal, ADSL mode, Speed: ADSL

  Speed: 832kbps, Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:b1:7e:85:84:ff
  Last flapped   : 2009-10-28 22:18:50 PDT (00:37:35 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  ADSL alarms    : None
  ADSL defects   : None
  ADSL status:

```



```

Modem status : Showtime (Itu-dmt)
DSL mode      :      Auto      Annex A
Last fail code: None
Subfunction   : 0x00
Seconds in showtime : 2253

Logical interface at-1/0/0.0 (Index 71) (SNMP ifIndex 523)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: ATM-PPP-LLC
  Input packets : 36
  Output packets: 35
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 12 (00:00:07 ago), Output: 13 (00:00:05 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: 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: 1486
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
  Destination: 100.100.100.6, Local: 100.100.100.10
VCI 2.119
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 36
  Output packets: 35

Logical interface at-1/0/0.32767 (Index 72) (SNMP ifIndex 525)
  Flags: Point-To-Multipoint No-Multicast SNMP-Traps 0x0 Encapsulation: ATM-VCMUX

  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
VCI 2.4
  Flags: Active
  Total down time: 0 sec, Last down: Never
  Input packets : 0
  Output packets: 0

user@host> 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    100.100.100.10  --> 100.100.100.6
at-1/0/0.32767  up    up

[edit]
user@host# run ping 100.100.100.6 count 100 rapid
PING 100.100.100.6 (100.100.100.6): 56 data bytes

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 100.100.100.6 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 7.881/9.046/15.136/1.697 ms

```

Verifying PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and PAP Authentication

Purpose Verify the interface status and end-to-end data path connectivity

Action From operational mode, enter the **show interfaces pp0** and **show interfaces at-1/0/0 terse** commands.

```

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 526)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 63,
    Session AC name: belur, Remote MAC address: 00:90:1a:41:03:c5,
    Configured AC name: None, Service name: None,
    Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
    Underlying interface: at-1/0/0.0 (Index 71)
    Input packets : 464
    Output packets: 241
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1 (00:39:51 ago), Output: 225 (00:00:08 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: 1456
    Flags: Negotiate-Address
    Addresses, Flags: Kernel Is-Preferred Is-Primary
    Destination: 12.12.12.1, Local: 12.12.12.15

user@host> show interfaces at-1/0/0 terse
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			
at-1/0/0.32767	up	up			

```

[edit]
user@host# run show interfaces pp0 terse

```

Interface	Admin	Link	Proto	Local	Remote
pp0	up	up			
pp0.0	up	up	inet	12.12.12.15	--> 12.12.12.1

```

[edit]
user@host# run ping 12.12.12.1 count 100 rapid

```

```
PING 12.12.12.1 (12.12.12.1): 56 data bytes
```

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 12.12.12.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 9.369/10.590/16.716/1.660 ms
```

Verifying PPPoE over ATM on the at-1/0/0 Interface with Negotiated IP and CHAP Authentication

Purpose Verify the interface status and end-to-end data path connectivity

Action From operational mode, enter the **show interfaces pp0** and **show interfaces at-1/0/0 terse** commands.

```
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 70) (SNMP ifIndex 526)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 64,
    Session AC name: belur, Remote MAC address: 00:90:1a:41:03:c5,
    Configured AC name: None, Service name: None,
    Auto-reconnect timeout: 120 seconds, Idle timeout: Never,
    Underlying interface: at-1/0/0.0 (Index 71)
    Input packets : 14
    Output packets: 13
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 0 (never), Output: 7 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mp1s:
Not-configured
  CHAP state: Success
  PAP state: Closed
  Security: Zone: Null
  Protocol inet, MTU: 1456
  Flags: Negotiate-Address
  Addresses, Flags: Kernel Is-Preferred Is-Primary
    Destination: 12.12.12.1, Local: 12.12.12.16
```

```
user@host> 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
at-1/0/0.32767 up    up

[edit]
```

```

user@host# run show interfaces pp0 terse
Interface           Admin Link Proto   Local           Remote
pp0                  up    up
pp0.0                up    up   inet    12.12.12.16     --> 12.12.12.1

[edit]
user@host# run ping 12.12.12.1 count 1000 rapid

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 12.12.12.1 ping statistics ---
1000 packets transmitted, 1000 packets received, 0% packet loss
round-trip min/avg/max/stddev = 8.748/10.461/21.386/1.915 ms

[edit]
user@host#

```

Related Documentation

- [VDSL2 Interface Technology Overview on page 181](#)
- [Example: Configuring VDSL2 Interfaces \(Basic\) on page 220](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 227](#)

Example: Configuring VDSL2 Interfaces (Basic)

This example shows how to configure the VDSL2 interfaces for SRX110, SRX210, SRX220, SRX240, SRX320, and SRX340 devices. (Platform support depends on the Junos OS release in your installation.)

- [Requirements on page 220](#)
- [Overview on page 220](#)
- [Configuration on page 220](#)
- [Verifying the Configuration on page 223](#)

Requirements

Before you begin:

- Establish basic connectivity. See the *Quick Start Guide* for your device for factory default settings.
- Configure network interfaces as necessary. See [“Example: Creating an Ethernet Interface” on page 265](#).

Overview

In this example, you create a VDSL2 interface called **pt-1/0/0**, specify the type of encapsulation, and set the VDSL2 profile to auto.

Configuration

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 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 encapsulation ppp-over-ether
set interfaces pt-1/0/0 unit 0 family inet dhcp
set interfaces pt-1/0/0 unit 0 vlan-id 100
```

Step-by-Step Procedure To configure the VDSL2 interfaces for the SRX110, SRX210, SRX240, SRX320, and SRX340 devices and enable VLAN tagging:

1. Create an interface.

```
[edit]
user@host# edit interfaces pt-1/0/0
```
2. Set the type of VDSL2 profile.

```
[edit interfaces pt-1/0/0]
user@host# set vdsl-options vdsl-profile auto
```
3. Specify the logical unit to connect to this 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
```
5. To enable the DHCP client on the interface.

```
[edit interfaces pt-1/0/0]
user@host# set unit 0 family inet dhcp
```
6. Specify the type of encapsulation on the VDSL2 logical interface.

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



NOTE: The VDSL2 interface supports PPPoE. You can also set no encapsulation for the VDSL2 interface.



NOTE: To configure VLAN tagging, continue the configuration with the next step.

7. To enable VLAN tagging on the pt interface.

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

8. Specify the value of the VLAN ID to be configured.

```
[edit interfaces pt-1/0/0]
user@host# set interface pt-1/0/0 unit 0 vlan-id 100
```



NOTE: This feature is supported only on the pt interface, and the range of VLANs that can be configured is 0 to 4093.

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-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 pt-1/0/0
vdsl-options {
vdsl-profile auto;
}
unit 0 {
encapsulation ppp-over-ether;
Family inet {
address 100.100.100.1/24;
dhcp;
}
```



NOTE: When VLAN tagging is configured, the intended output is:

```
[edit]
user@host# show interfaces pt-1/0/0
vlan-tagging;
vdsl-options {
vdsl-profile auto;
}
unit 0 {
encapsulation ppp-over-ether;
vlan-id 100;
Family inet {
address 100.100.100.1/24;
dhcp;
}
```

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

Verifying the Configuration

Confirm that the configuration is working properly.

- [Displaying the Configuration for VDSL2 Interface \(When Connected to the DSLAM Operating in Annex A Mode\)](#) on page 223
- [Displaying the Configuration for VDSL2 Interface \(When Connected to the DSLAM Operating in Annex B Mode\)](#) on page 225

Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex A Mode)

Purpose Verify the command output.

Action From operational mode, enter the **show interfaces pt-1/0/0** command.

```
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
VDSL Chipset Information: VTU-R VTU-C
Vendor Country : 0xb5 0xb5
Vendor ID : BDCM BDCM
Vendor Specific: 0x9385 0x9385
VDSL Statistics: VTU-R VTU-C
Attenuation (dB) : 0.0 0.0
Capacity used (%) : 0 0
Noise margin (dB) : 20.0 20.0
Output power (dBm) : 6.0 12.0

          Interleave Fast Interleave Fast
Bit rate (kbps) :    100004      0    45440      0
CRC :              0          0      0          0
FEC :              0          0      0          0
HEC :              0          0      0          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      43168000    95          0      low
none
3 network-control 5        2272000     5          0      low
none
Logical interface pt-1/0/0.0 (Index 71) (SNMP ifIndex 525) (Generation 136)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes : 23789064
Output bytes : 10866024
Input packets: 16052
Output packets: 7332
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 23789064 97070256 bps
Output bytes : 10866024 43334088 bps
Input packets: 16052 8187 pps
Output packets: 7332 3655 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: 1482, Generation: 169, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary

```

```

Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 158

```

The output shows a summary of VDSL2 interface. Verify the following information:

- Status of interface pt-1/0/0 is displayed as Physical link is Up.
- Modem status is displayed as Showtime (Profile-17a).
- Time in seconds during which the interface stayed up is displayed as Seconds in showtime.
- Annex A indicates VDSL profile of the DSLAM connected at other end.

Displaying the Configuration for VDSL2 Interface (When Connected to the DSLAM Operating in Annex B Mode)

Purpose Verify the command output.

Action From operational mode, enter the **show interfaces pt-1/0/0** command.

```

Physical interface: pt-1/0/0, Enabled, Physical link is Up
Interface index: 148, SNMP ifIndex: 536, Generation: 238
Type: PTM, Link-level type: Ethernet, MTU: 1514, VDSL mode, Speed: 45439kbps
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: 00:1f:12:e4:df:20
Last flapped   : 2011-05-13 07:34:33 PDT (00:46:33 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
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: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,
Resource errors: 0
VDSL alarms : None
VDSL defects : None
VDSL media:
  Seconds      Count  State
  LOF          177    0 OK
  LOS          177    0 OK
  LOM           0     0 OK
  LOP           0     0 OK
  LOCDI         0     0 OK
  LOCDNI        177    0 OK
VDSL status:
  Modem status : Showtime (Profile-17a)
  VDSL profile : Auto Annex B
  Last fail code: None
  Subfunction : 0x00
Seconds in showtime: 2794 VDSL Chipset Information: VTU-R
  VTU-C
  Vendor Country : 0xb5 0xb5
  Vendor ID : BDCM BDCM
  Vendor Specific: 0x9385 0x9395
VDSL Statistics: VTU-R VTU-C
  Attenuation (dB) : 0.0 0.0
  Capacity used (%) : 0 0
  Noise margin (dB) : 18.5 9.5
  Output power (dBm) : 14.5 3.0

  Interleave Fast Interleave Fast
  Bit rate (kbps) : 100015 0 45439 0
  CRC : 0 0 0 0
  FEC : 0 0 0 0
  HEC : 0 0 0 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 43167050 95 0 low
none
  3 network-control 5 2271950 5 0 low
none

```

The output shows a summary of the VDSL2 interface. Verify the following information:

- Status of interface pt-1/0/0 is displayed as Physical link is Up.
- Modem status is displayed as Showtime (Profile-17a).

- Time in seconds during which the interface stayed up is displayed as Seconds in showtime.
- Annex B indicates the VDSL profile of the DSLAM connected at other end.

Related Documentation

- [Understanding Interfaces on page 3](#)
- [VDSL2 Interface Technology Overview on page 181](#)
- [Example: Configuring VDSL2 Interfaces \(Detail\) on page 227](#)
- [Example: Configuring VDSL2 Interfaces in ADSL Mode \(Detail\) on page 193](#)

Example: Configuring VDSL2 Interfaces (Detail)

This example shows how to configure VDSL2 interfaces on SRX Series Services Gateways.

This example uses VDSL2 Mini-PIM installed on SRX210 and SRX320 devices. The information is also applicable to SRX110 (integrated VDSL2), SRX220, SRX240, and SRX320 devices (with VDSL2 Mini-PIMs). (Platform support depends on the Junos OS release in your installation.)

- [Requirements on page 227](#)
- [Overview on page 227](#)
- [Configuration on page 228](#)
- [Verification on page 241](#)

Requirements

Before you begin:

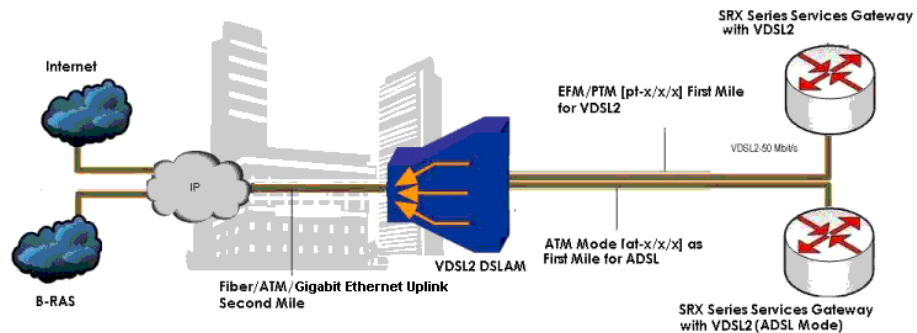
- Install Junos OS Release 10.1 or later on the SRX Series devices.
- Establish basic connectivity and set up and perform initial configuration. See the *Quick Start Guide* for your device for factory default settings.
- Install the VDSL2 Mini-PIM on the SRX210 or SRX320 device chassis.
- Connect the SRX210 or SRX320 device to a DSLAM.
- On VDSL2 Mini-PIMs, by default the **pt-1/0/0** interface is created when there is no configuration already created for either the **pt-1/0/0** or the **at-1/0/0** interface. You can switch to ADSL mode by just configuring **at-1/0/0**. If the configurations are already created for **pt-1/0/0** or **at-1/0/0**, then you need to 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**.

Overview

This example uses SRX210 or SRX320 devices. The information is also applicable to SRX240 and SRX340 devices.

Figure 17 on page 228 shows typical SRX Series devices with VDSL2 Mini-PIM network connections.

Figure 17: SRX Series Device with VDSL2 Mini-PIMs in an End-to-End Deployment Scenario



In this example, you begin a new configuration on a VDSL2 Mini-PIM. You first deactivate previous interfaces and delete any old configuration from the device. Then you set the interfaces with the VDSL profile and the Layer 3 configuration for the end-to-end data path.

You then configure the PPPoE on the pt-1/0/0 interface with a static IP address or CHAP authentication. You configure PPPoE on the pt-1/0/0 interface with unnumbered IP address (PAP authentication or CHAP authentication).

Finally, you configure PPPoE on the pt-1/0/0 interface with negotiated IP address (PAP authentication or CHAP authentication).

Configuration

- [Beginning a New Configuration on a VDSL2 Mini-PIM on page 229](#)
- [Configuring the VDSL2 Mini-PIM for End-to-End Data Path on page 229](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address on page 230](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address \(CHAP Authentication\) on page 232](#)
- [Configuring PPPoE on the pt-x/x/x Interface with Unnumbered IP \(PAP Authentication\) on page 234](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with Unnumbered IP \(CHAP Authentication\) on page 236](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP \(PAP Authentication\) on page 238](#)
- [Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP \(CHAP Authentication\) on page 239](#)

Beginning a New Configuration on a VDSL2 Mini-PIM

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.

```
[edit]
deactivate interface pt-1/0/0
deactivate interface at-1/0/0
delete interface pt-1/0/0
delete interface pp0
```

Step-by-Step Procedure 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. 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 chassis fpc** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# run 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 -----
```

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

Configuring the VDSL2 Mini-PIM for End-to-End Data Path

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 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
```

Step-by-Step Procedure

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 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. 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 pt-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 pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  family inet {
    address 11.11.11.1/24;
  }
}
```

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

Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address

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.

```
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
```



NOTE: To configure VLAN tagging while configuring PPPoE on the pt-1/0/0 interface with

- Static IP address
- Static IP address (CHAP authentication)
- Unnumbered IP address (PAP Authentication)
- Unnumbered IP address (CHAP Authentication)
- Negotiated IP address (PAP Authentication)
- Negotiated IP address (CHAP Authentication)

the following commands must be included at [edit] hierarchy level:

```
set interfaces pt-1/0/0 vlan-tagging
set interfaces pt-1/0/0 unit 0 vlan-id 100
```

Step-by-Step Procedure

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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0**, **show interfaces pt-1/0/0** and **show access profile pap_prof** 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 pp0
unit 0 {
  ppp-options {
    pap {
      access-profile pap_prof;
      local-name locky;
      local-password "$ABC123"; ## SECRET-DATA
    }
    passive;
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    address 10.1.1.6/24;
  }
}
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  encapsulation ppp-over-ether;
}
[edit]
user@host# show access profile pap_prof
authentication-order password;
client cuttack pap-password "$ABC123"; ## SECRET-DATA
```

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

Configuring PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication)

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.

```
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

```

Step-by-Step Procedure 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

```

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0** and **show interfaces pp0** 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 pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  encapsulation ppp-over-ether;
}
[edit]
user@host# show interfaces pp0
unit 0 {
  ppp-options {
    chap {
      default-chap-secret "$ABC123"; ## SECRET-DATA

```

```
local-name locky;
passive;
}
}
pppoe-options {
underlying-interface pt-1/0/0.0;
auto-reconnect 120;
client;
}
family inet {
address 10.1.1.6/24;
}
}
```

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

Configuring PPPoE on the pt-x/x/x Interface with Unnumbered IP (PAP Authentication)

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.

```
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
```

Step-by-Step Procedure

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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces lo0**, **show interfaces pt-1/0/0**, and **show interfaces pp0** 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 lo0
unit 0 {
family inet {
address 10.1.1.24/32;
}
}
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
vdsl-profile 17a;
}
unit 0 {
encapsulation ppp-over-ether;
}
[edit]
user@host# show interfaces pp0
unit 0 {
ppp-options {
pap {
access-profile pap_prof;
local-name locky;
local-password "$ABC123"; ## SECRET-DATA
passive;
}
}
}
```

```

    }
    pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
    }
    family inet {
    unnumbered-address lo0.0 destination 10.1.1.1;
    }
    }

```

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

Configuring PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication)

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.

```

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

```

Step-by-Step Procedure

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

```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0**, **show interfaces pt-1/0/0**, and **show interfaces lo0** 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 pp0
unit 0 {
  ppp-options {
    chap {
      default-chap-secret "$ABC123"; ## SECRET-DATA
      local-name locky;
      passive;
    }
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    unnumbered-address lo0.0 destination 10.1.1.1;
  }
}
[edit]
user@host# show interfaces pt-1/0/0
vdsl-options {
  vdsl-profile 17a;
}
unit 0 {
  encapsulation ppp-over-ether;
}
[edit]
user@host# show interfaces lo0
unit 0 {
  family inet {
    address 10.1.1.24/32;
  }
}

```

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

Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication)

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.

```
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
```

Step-by-Step Procedure 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

Results From configuration mode, confirm your configuration by entering the **show interfaces pt-1/0/0**, **show interfaces pp0**, and **show access profile my_prf** 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 pt-1/0/0
vdsl-options {
vdsl-profile 17a;
}
unit 0 {
encapsulation ppp-over-ether;
}
[edit]
user@host# show interfaces pp0
unit 0 {
ppp-options {
pap {
access-profile my_prf;
local-name purple;
local-password "$ABC123"; ## SECRET-DATA
passive;
}
}
pppoe-options {
underlying-interface pt-1/0/0.0;
auto-reconnect 120;
client;
}
family inet {
negotiate-address;
}
}
[edit]
user@host# show access profile my_prf
authentication-order password;
```

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

Configuring PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication)

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.

```
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
```

Step-by-Step Procedure 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
```

Results From configuration mode, confirm your configuration by entering the **show interfaces pp0** and **show interfaces pt-1/0/0** 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 pp0
unit 0 {
  ppp-options {
    chap {
      default-chap-secret "$ABC123"; ## SECRET-DATA
      local-name purple;
      passive;
    }
  }
  pppoe-options {
    underlying-interface pt-1/0/0.0;
    auto-reconnect 120;
    client;
  }
  family inet {
    negotiate-address;
  }
}
[edit]
user@host# show interfaces pt-1/0/0
```



```

vdsl-options {
vdsl-profile 17a;
}
unit 0 {
encapsulation ppp-over-ether;
}

```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the Configuration on page 241](#)
- [Verifying the VDSL2 Mini-PIM for End-to-End Data Path on page 244](#)
- [Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address on page 247](#)
- [Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address \(CHAP Authentication\) on page 248](#)
- [Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP \(PAP Authentication\) on page 249](#)
- [Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP \(CHAP Authentication\) on page 250](#)
- [Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP \(PAP Authentication\) on page 251](#)
- [Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP \(CHAP Authentication\) on page 252](#)

Verifying the Configuration

Purpose Verify the FPC status and the command output.

Action 1. Verify the FPC status by entering the **show chassis fpc** command. The output should display FPC status as online.

```

user@host# run show chassis fpc
Temp CPU Utilization (%) Memory Utilization
(%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online ----- CPU 1ess FPC -----
1 Online ----- CPU 1ess FPC -----

```



NOTE: The VDSL2 Mini-PIM is installed in the first slot of the SRX320 device chassis; therefore, the FPC used here is fpc 1. For SRX340 devices, the FPC used will be fpc 1, fpc 2, fpc 3, or fpc 4.

2. Enter **run show interface pt-1/0/0** and verify the following information in the command output:

- Status of interface pt-1/0/0 is displayed as physical link is up.
- Modem status is displayed as Showtime (Profile-17a).
- Time in seconds during which the interface stayed up is displayed as Seconds in Showtime.
- VDSL profile of DSLAM is displayed as Auto Annex A.

```
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
VDSL Chipset Information: VTU-R VTU-C
  Vendor Country : 0xb5 0xb5
  Vendor ID : BDCM BDCM
```

```

Vendor Specific: 0x9385 0x9385
VDSL Statistics: VTU-R VTU-C
Attenuation (dB) : 0.0 0.0
Capacity used (%) : 0 0
Noise margin (dB) : 20.0 20.0
Output power (dBm) : 6.0 12.0

                Interleave Fast Interleave Fast
Bit rate (kbps) :    100004      0    45440      0
CRC :              0          0    0          0
FEC :              0          0    0          0
HEC :              0          0    0          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      43168000    95          0      low
none
3 network-control 5        2272000    5          0      low
none
Logical interface pt-1/0/0.0 (Index 71) (SNMP ifIndex 525) (Generation 136)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes : 23789064
  Output bytes : 10866024
  Input packets: 16052
  Output packets: 7332
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 23789064 97070256 bps
  Output bytes : 10866024 43334088 bps
  Input packets: 16052 8187 pps
  Output packets: 7332 3655 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: 1482, Generation: 169, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 158

```

Verifying the VDSL2 Mini-PIM for End-to-End Data Path

Purpose Verify the interface status and check traffic statistics.

- Action** 1. Verify interface status by using 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 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
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
- 11.11.11.2 ping statistics ---
1000 packets transmitted, 1000 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.109/17.711/28.591/2.026 ms

```

2. Verify the VDSL2 interface configuration and check the traffic statistics.

```

user@host# run show interfaces pt-1/0/0 extensive
Physical interface: pt-1/0/0, Enabled, Physical link is Up
Interface index: 146, SNMP ifIndex: 524, Generation: 197
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   : 2009-10-28 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
  Subfunction    : 0x00
  Seconds in showtime : 723
VDSL Chipset Information:
  Vendor Country : VTU-R          VTU-C
  Vendor ID      : 0xb5          0xb5
  Vendor Specific: BDCM          BDCM
  Vendor Specific: 0x9385        0x9385
VDSL Statistics:
  Attenuation (dB) : VTU-R          VTU-C
  Capacity used (%) : 0.0          0.0
  Noise margin (dB) : 0          0
  Output power (dBm) : 16.0        20.0
  Output power (dBm) : 5.0         13.0

Interleave      Fast  Interleave      Fast
Bit rate (kbps) : 100004      0      45440
0
CRC              : 0          0          0
0
FEC              : 0          0          0
0
HEC              : 0          0          0
0
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue
Limit          Bandwidth          Buffer Priority
              %          bps          %          usec

```

0 best-effort	95	43168000	95	0	low
none					
3 network-control	5	2272000	5	0	low
none					

Logical interface pt-1/0/0.0 (Index 72) (SNMP ifIndex 521) (Generation 158)

Flags: SNMP-Traps Encapsulation: ENET2

Traffic statistics:

Input bytes :	84000
Output bytes :	98000
Input packets:	1000
Output packets:	1000

Local statistics:

Input bytes :	84000
Output bytes :	98000
Input packets:	1000
Output packets:	1000

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: 1482, Generation: 169, Route table: 0

Flags: None

Addresses, Flags: Is-Preferred Is-Primary

Destination: 11.11.11/24, Local: 11.11.11.1, Broadcast: 11.11.11.255,
Generation: 189

Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address

Purpose Verify the interface output and the end-to-end data path.

Action 1. Verify the interface output.

```
user@host# run 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
  NCP state: inet: Down, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Closed
  PAP state: Closed
  Security: Zone: Null
  Protocol inet, MTU: 1492
  Flags: Protocol-Down
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.1.1/24, Local: 10.1.1.6
```

2. Verify the end-to-end data path on the interface.

```
user@host# run 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
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
- 10.1.1.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 14.669/15.649/21.655/1.740 ms

```

Verifying PPPoE on the pt-1/0/0 Interface with a Static IP Address (CHAP Authentication)

Purpose Verify the interface status and check the end-to-end data path connectivity.

Action 1. Verify the interface status.

```

user@host# run 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 70) (SNMP ifIndex 522)
  Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 31,
    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 : 12
  Output packets: 10
  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: Success
  PAP state: Closed
  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

```

2. Verify the interface and check the end-to-end data path connectivity.

```

user@host# run 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

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 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

```

Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (PAP Authentication)

Purpose Verify the interface status and the end-to-end data path testing.

Action 1. Verify the interface status.

```

user@host# run 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

```

2. Verify the end-to-end data path testing.

```

user@host# run 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    10.1.1.24   --> 10.1.1.1

[edit]
user@host# run 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 = 14.584/15.503/21.204/1.528 ms

```

Verifying PPPoE on the pt-1/0/0 Interface with Unnumbered IP (CHAP Authentication)

Purpose Verify the interface status and end-to-end data path testing on the PPPoE interface.

Action 1. Verify the interface status.

```

user@host# run 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 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

2. Verify the end-to-end data path testing on the PPPoE interface.

```
user@host# run 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    10.1.1.24   --> 10.1.1.1

[edit]
user@host# run 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 = 14.585/16.025/22.354/2.019 ms
```

Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (PAP Authentication)

Purpose Verify the PPPoE interface status and the end-to-end data path connectivity.

- Action** 1. Verify the PPPoE interface status.

```
user@host# run 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: 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

```

2. Verify the end-to-end data path connectivity.

```

user@host# run 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    12.12.12.11  --> 12.12.12.1

[edit]
user@host# run ping 12.12.12.1 count 100 rapid
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 12.12.12.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.223/17.692/24.359/2.292 ms

```

Verifying PPPoE on the pt-1/0/0 Interface with Negotiated IP (CHAP Authentication)

Purpose Verify the interface status and the end-to-end data path connectivity.

- Action** 1. Verifying the interface status.

```

user@host# run 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 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

```

2. Verify the end-to-end data path connectivity.

```

user@host# run 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    12.12.12.12  --> 12.12.12.1

[edit]
user@host# run ping 12.12.12.1 count 100 rapid
PING 12.12.12.1 (12.12.12.1): 56 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 12.12.12.1 ping statistics ---
100 packets transmitted, 100 packets received, 0% packet loss
round-trip min/avg/max/stddev = 16.168/17.452/23.299/2.016 ms

```

- Related Documentation**
- [VDSL2 Interface Technology Overview on page 181](#)
 - [Example: Configuring VDSL2 Interfaces \(Basic\) on page 220](#)
 - [Example: Configuring VDSL2 Interfaces in ADSL Mode \(Detail\) on page 193](#)

Upgrading the VDSL PIC Firmware

Starting with Junos OS Release 15.1x49-D50, you can upgrade the VDSL PIC firmware on SRX Series devices. This topic shows how to perform the upgrade.

Before you begin:

Check the current firmware version of the VDSL PIC.

```

user@host> show system firmware
Part Type Tag Current Available Status version
FPC 1
PIC 0 VDSLBCM 10 2.10.0 OK
Routing Engine 0 RE BIOS 0 2.0 OK
Routing Engine 0 RE BIOS Backup 1 2.0 OK
Routing Engine 0 RE FPGA 14 1.0.0 OK

```

This section describes the step-by-step procedure to upgrade VDSL PIC firmware.

1. Mount or copy the firmware package to the SRX Series device.

If the file has been obtained from JTAC, use FTP or SCP to load the firmware file on the device. Save the file in the `/var/tmp` directory.

2. Upgrade the firmware on the SRX Series device.

To install the firmware package on the device and make it available for upgrading, use the following command:

```
user@host> request system software add no-copy no-validate  
jfirmware-srxsme-11.4R2.7-signed.tgz
```

3. To check if the firmware package is available on the SRX Series device, use the following command:

```
user@host> show version  
  
Hostname: user  
  
Model: srx210h  
  
JUNOS Software Release [12.1I20120123_0941]  
  
JUNOS Firmware Software Suite [11.4R2.7]
```

4. To verify the VDSL PIM slot, use the following command:

```
user@host> show chassis hardware
```

5. To initiate a firmware upgrade, use the following command:

```
user@host> request system firmware upgrade pic fpc-slot <no.> pic-slot 0 tag 10
```

6. To check the status of the upgraded firmware, use the following command:

```
user@host> show system firmware  
Part Type Tag Current Available Status version  
FPC 1  
PIC 0 VDSLBCM 10 2.10.0 2.11.0  
Routing Engine 0 RE BIOS 0 2.0 OK  
Routing Engine 0 RE BIOS Backup 1 2.0 OK  
Routing Engine 0 RE FPGA 14 203.0.113.45.0.0 OK
```

7. To enable the upgraded firmware, restart the FPC slot in which the VDSL PIM is installed.

```
user@host> restart fpc <no.>  
  
FPC 1 restarted
```

8. To verify the firmware upgrade is complete, use the following command:

```
user@host> show system firmware
Part Type Tag Current Available Status version
FPC 1
PIC 0 VDSLBCM 10 2.11.0 2.11.0 OK
Routing Engine 0 RE BIOS 0 2.0 OK
Routing Engine 0 RE BIOS Backup 1 2.0 OK
Routing Engine 0 RE FPGA 14 203.0.113.45.0.0 OK
```

Release History Table

Release	Description
15.1X49-D50	Starting with Junos OS Release 15.51x49-D50, you can upgrade the VDSL PIC firmware on SRX Series devices.

PART 4

Configuring Ethernet Interfaces

- [Performing Initial Configuration on Ethernet Interfaces on page 259](#)
- [Configuring Aggregated Ethernet Interfaces on page 279](#)
- [Configuring Link Aggregation Control Protocol on page 295](#)
- [Configuring Gigabit Ethernet Physical Interface Modules on page 317](#)
- [Configuring Port Mirroring on page 351](#)
- [Configuring Ethernet OAM Link Fault Management on page 355](#)
- [Configuring Power over Ethernet on page 367](#)

CHAPTER 13

Performing Initial Configuration on Ethernet Interfaces

- [Understanding Ethernet Interfaces on page 259](#)
- [Understanding Static ARP Entries on Ethernet Interfaces on page 263](#)
- [Understanding Promiscuous Mode on Ethernet Interface on page 263](#)
- [Understanding Port Mirroring on SRX Devices on page 264](#)
- [Example: Creating an Ethernet Interface on page 265](#)
- [Example: Deleting an Ethernet Interface on page 266](#)
- [Example: Configuring Static ARP Entries on Ethernet Interfaces on page 267](#)
- [Enabling and Disabling Promiscuous Mode on Ethernet Interfaces \(CLI Procedure\) on page 270](#)
- [Example: Configuring Promiscuous Mode on the SRX5K-MPC on page 271](#)
- [Configuring Port Mirroring on SRX Devices on page 275](#)

Understanding Ethernet Interfaces

Ethernet is a Layer 2 technology that operates in a shared bus topology. Ethernet supports broadcast transmission, uses best-effort delivery, and has distributed access control. Ethernet is a point-to-multipoint technology.

In a shared bus topology, all devices connect to a single, shared physical link through which all data transmissions are sent. All traffic is broadcast so that all devices within the topology receive every transmission. The devices within a single Ethernet topology make up a broadcast domain.

Ethernet uses best-effort delivery to broadcast traffic. The physical hardware provides no information to the sender about whether the traffic was received. If the receiving host is offline, traffic to the host is lost. Although the Ethernet data link protocol does not inform the sender about lost packets, higher layer protocols such as TCP/IP might provide this type of notification.

This topic contains the following sections:

- [Ethernet Access Control and Transmission on page 260](#)
- [Collisions and Detection on page 260](#)

- [Collision Domains and LAN Segments on page 261](#)
- [Broadcast Domains on page 262](#)
- [Ethernet Frames on page 262](#)

Ethernet Access Control and Transmission

Ethernet's access control is distributed because Ethernet has no central mechanism that grants access to the physical medium within the network. Instead, Ethernet uses carrier-sense multiple access with collision detection (CSMA/CD). Because multiple devices on an Ethernet network can access the physical medium, or wire, simultaneously, each device must determine whether the physical medium is in use. Each host listens on the wire to determine if a message is being transmitted. If it detects no transmission, the host begins transmitting its own data.

The length of each transmission is determined by fixed Ethernet packet sizes. By fixing the length of each transmission and enforcing a minimum idle time between transmissions, Ethernet ensures that no pair of communicating devices on the network can monopolize the wire and block others from sending and receiving traffic.

Collisions and Detection

When a device on an Ethernet network begins transmitting data, the data takes a finite amount of time to reach all hosts on the network. Because of this delay, or latency, in transmitting traffic, a device might detect an idle state on the wire just as another device initially begins its transmission. As a result, two devices might send traffic across a single wire at the same time. When the two electrical signals collide, they become scrambled so that both transmissions are effectively lost.

Collision Detection

To handle collisions, Ethernet devices monitor the link while they are transmitting data. The monitoring process is known as collision detection. If a device detects a foreign signal while it is transmitting, it terminates the transmission and attempts to transmit again only after detecting an idle state on the wire. Collisions continue to occur if two colliding devices both wait the same amount of time before retransmitting. To avoid this condition, Ethernet devices use a binary exponential backoff algorithm.

Backoff Algorithm

With the binary exponential backoff algorithm, each device that sends a colliding transmission randomly selects a value within a range. The value represents the number of transmission times that the device must wait before retransmitting its data. If another collision occurs, the range of values is doubled and retransmission takes place again. Each time a collision occurs, the range of values doubles, to reduce the likelihood that two hosts on the same network can select the same retransmission time.

[Table 34 on page 260](#) shows collision rounds up to round 10.

Table 34: Collision Backoff Algorithm Rounds

Round	Size of Set	Elements in the Set
1	2	{0,1}

Table 34: Collision Backoff Algorithm Rounds (continued)

Round	Size of Set	Elements in the Set
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}

Collision Domains and LAN Segments

Collisions are confined to a physical wire over which data is broadcast. Because the physical wires are subject to signal collisions, individual LAN segments are known as *collision domains*. Although the physical limitations on the length of an Ethernet cable restrict the length of a LAN segment, multiple collision domains can be interconnected by repeaters, bridges, and switches.

Repeaters

Repeaters are electronic devices that act on analog signals. Repeaters relay all electronic signals from one wire to another. A single repeater can double the distance between two devices on an Ethernet network. However, the Ethernet specification restricts the number of repeaters between any two devices on an Ethernet network to two, because collision detection with latencies increases in complexity as the wire length and number of repeaters increase.

Bridges and Switches

Bridges and switches combine LAN segments into a single Ethernet network by using multiple ports to connect the physical wires in each segment. Although bridges and switches are fundamentally the same, bridges generally provide more management and more interface ports. As Ethernet packets flow through a bridge, the bridge tracks the source MAC address of the packets and stores the addresses and their associated input ports in an interface table. As it receives subsequent packets, the bridge examines its interface table and takes one of the following actions:

- If the destination address does not match an address in the interface table, the bridge transmits the packet to all hosts on the network using the Ethernet broadcast address.

- If the destination address maps to the port through which the packet was received, the bridge or switch discards the packet. Because the other devices on the LAN segment also received 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.

Broadcast Domains

The combination of all the LAN segments within an Ethernet network is called a *broadcast domain*. In the absence of any signaling devices such as a repeater, bridge, or switch, the broadcast domain is simply the physical wire that makes up the connections in the network. If a bridge or switch is used, the broadcast domain consists of the entire LAN.

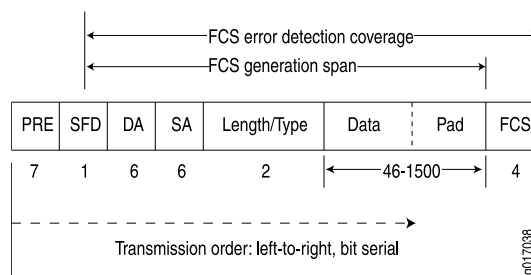


NOTE: On SRX300, SRX320, SRX340, SRX345, and SRX550HM devices, the subnet directed broadcast feature is not supported.

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 18 on page 262](#) shows the Ethernet frame format.

Figure 18: 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.



NOTE: 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.)

Related Documentation

- [Understanding Interfaces on page 3](#)
- [Example: Creating an Ethernet Interface on page 265](#)
- [Example: Deleting an Ethernet Interface on page 266](#)
- [Understanding Static ARP Entries on Ethernet Interfaces on page 263](#)
- [Understanding Promiscuous Mode on Ethernet Interface on page 263](#)

Understanding Static ARP Entries on Ethernet Interfaces

By default, the device responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the device to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Example: Configuring Static ARP Entries on Ethernet Interfaces on page 267](#)

Understanding Promiscuous Mode on Ethernet Interface

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. 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, 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.

Understanding Promiscuous Mode on the SRX5K-MPC

The 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).

When promiscuous mode is enabled on a Layer 3 Ethernet interface, all packets received on the interface are sent to the central point or to the Services Processing Unit (SPU) regardless of the destination MAC address of the packet.

By default, an interface enables MAC filtering. You can configure promiscuous mode on the interface to disable MAC filtering. When you delete the promiscuous mode configuration, the interface will perform MAC filtering again.

You can change the MAC address of an interface even when the interface is operating in promiscuous mode. When the interface is operating in normal mode again, the MAC filtering function on the IOC uses the new MAC address to filter the packets.

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, 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.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Enabling and Disabling Promiscuous Mode on Ethernet Interfaces \(CLI Procedure\) on page 270](#)
- [Example: Configuring Promiscuous Mode on the SRX5K-MPC on page 271](#)

Understanding Port Mirroring on SRX Devices

Port mirroring copies packets entering or exiting a port and sends the copies to a local interface for monitoring. Port mirroring is used to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Port mirroring is used to send a copy of all the packets or only the sampled packets seen on a port to a network monitoring connection. You can mirror the packets either on the incoming port (ingress port mirroring) or the outgoing port (egress port mirroring).



NOTE: Port mirroring is supported only on the SRX devices with the following I/O cards:

- SRX1K-SYSIO-GE
- SRX1K-SYSIO-XGE
- SRX3K-SFB-12GE
- SRX3K-2XGE-XFP
- SRX5K-FPC-IOC Flex I/O

On SRX devices, all packets passing through the **mirrored** port are copied and sent to the specified **mirror-to** port. These ports must be on the same Broadcom chipset in the I/O cards.



NOTE: On SRX devices, port mirroring works on physical interfaces only.

Related Documentation

- [Configuring Port Mirroring on SRX Devices on page 275](#)

Example: Creating an Ethernet Interface

This example shows how to create an Ethernet interface.

- [Requirements on page 265](#)
- [Overview on page 265](#)
- [Configuration on page 265](#)

Requirements

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

Overview

In this example, you create the ge-1/0/0 Ethernet interface and set the logical interface to 0. The logical unit number can range from 0 to 16,384. You can also add values for properties that you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuration

Step-by-Step Procedure

To configure an Ethernet interface:

1. Create the Ethernet interface and set the logical interface.

[edit]

user@host# edit interfaces ge-1/0/0 unit 0

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

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

Verification

Purpose Verify if the configuration is working properly after creating the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Example: Deleting an Ethernet Interface on page 266](#)

Example: Deleting an Ethernet Interface

This example shows how to delete an Ethernet interface.

- [Requirements on page 266](#)
- [Overview on page 266](#)
- [Configuration on page 266](#)

Requirements

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

Overview

In this example, you delete the ge-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 J-Web pages.

Configuration

Step-by-Step Procedure

To delete an Ethernet interface:

1. Specify the interface you want to delete.

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

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

```
[edit]
```

```
user@host# commit
```

Verification

Purpose Verify if the configuration is working properly after deleting the interface.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Example: Creating an Ethernet Interface on page 265](#)

Example: Configuring Static ARP Entries on Ethernet Interfaces

- [Requirements on page 267](#)
- [Overview on page 267](#)
- [Configuration on page 267](#)
- [Verification on page 268](#)

Requirements

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

Overview

In this example, you configure a static ARP entry on the logical unit 0 of the ge-0/0/3 Gigabit Ethernet interface. The entry consists of the interface's IP address (10.1.1.1/24) and the corresponding MAC address of a node on the same Ethernet subnet (00:ff:85:7f:78:03). The example also configures the device to reply to ARP requests from the node using the publish option.

Configuration

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-0/0/3 unit 0 family inet address 10.1.1.1/24 arp 10.1.1.3 mac  
00:ff:85:7f:78:03  
set interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24 arp 10.1.1.3 publish
```

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 static ARP entry on an Ethernet interface:

1. Create the Gigabit Ethernet interface.

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

2. Configure a static ARP entry.

```
[edit interfaces ge-0/0/3]
user@host# edit unit 0 family inet address 10.1.1.1/24
```

3. Set the IP address of the subnet node and the corresponding MAC address.

```
[edit interfaces ge-0/0/3 unit 0 family inet address 10.1.1.1/24]
user@host# set arp 10.1.1.3 mac 00:ff:85:7f:78:03 publish
```

Results From configuration mode, confirm your configuration by entering the **show interfaces ge-0/0/3** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces ge-0/0/3
unit 0 {
  family inet {
    address 10.1.1.1/24 {
      arp 10.1.1.3 mac 00:ff:85:7f:78:03 publish;
    }
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying Static ARP Configurations on page 268](#)
- [Verifying the Link State of All Interfaces on page 269](#)
- [Verifying Interface Properties on page 269](#)

Verifying Static ARP Configurations

Purpose Verify the IP address and MAC (hardware) address of the node.

Action From operational mode, enter the **show interfaces ge-0/0/3** command.

Verifying the Link State of All Interfaces

Purpose Verify that all interfaces on the device are operational using the ping tool on each peer address in the network.

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 10.1.1.3 : 56 data bytes
64 bytes from 10.1.1.3: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.1.1.3: 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 operational mode, enter the **show interfaces detail** command.

```
user@host> show interfaces detail
Physical interface: ge-0/0/3, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 27, Generation: 17
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  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: 00:90:69:87:44:9d, Hardware address: 00:90:69:87:44:9d
  Last flapped   : 2004-08-25 15:42:30 PDT (4w5d 22:49 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
  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

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 ge-0/0/3] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the Interfaces > ge-0/0/3 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 expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics ge-0/0/3** command.

**Related
Documentation**

- [Understanding Static ARP Entries on Ethernet Interfaces on page 263](#)

Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure)

To enable promiscuous mode on an interface:

```
user@host# set interfaces interface-name promiscuous-mode
```

To disable promiscuous mode on an interface:

```
user@host# delete interfaces interface-name promiscuous-mode
```

**Related
Documentation**

- [Understanding Promiscuous Mode on Ethernet Interface on page 263](#)
- [Understanding Ethernet Interfaces on page 259](#)

Example: Configuring Promiscuous Mode on the SRX5K-MPC

This example shows how to configure promiscuous mode on an SRX5K-MPC interface in an SRX5600 to disable MAC address filtering.

- [Requirements on page 271](#)
- [Overview on page 271](#)
- [Configuration on page 271](#)
- [Verification on page 272](#)

Requirements

This example uses the following hardware and software components:

- An SRX5600 with an SRX5K-MPC that includes a 100-Gigabit Ethernet CFP transceiver
- Junos OS Release 12.1X47-D10 or later

No special configuration beyond device initialization is required before configuring this feature.

Overview

By default, the interfaces on an SRX5K-MPC have MAC address filtering enabled. In this example, you configure promiscuous mode on an interface to disable MAC address filtering. Then you delete promiscuous mode to reenable MAC address filtering on the interface.

Configuration

Configuring Promiscuous Mode on an 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 et-4/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces et-4/0/0 promiscuous-mode
```

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 promiscuous mode:

1. Configure the ingress interface.

```
[edit interfaces]
user@host# set et-4/0/0 unit 0 family inet address 10.1.1.1/24
```
2. Enable promiscuous mode on the interface.

```
[edit interfaces]
user@host# set et-4/0/0 promiscuous-mode
```

Results From configuration mode, confirm your configuration by entering the **show** 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
et-4/0/0 {
  promiscuous-mode;
  unit 0 {
    family inet {
      address 10.1.1.1/24;
    }
  }
}
```

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

Disabling Promiscuous Mode on an 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.

```
user@host# delete interfaces et-4/0/0 promiscuous-mode
```

Step-by-Step Procedure To disable promiscuous mode:

1. Disable promiscuous mode on the interface.

```
[edit]
user@host# delete interfaces et-4/0/0 promiscuous-mode
```

Verification

Confirm that the configuration is working properly.

- [Verifying That Promiscuous Mode Is Enabled on the SRX5K-MPC on page 272](#)
- [Verifying the Status of Promiscuous Mode on page 273](#)
- [Verifying That Promiscuous Mode Is Disabled on page 274](#)

Verifying That Promiscuous Mode Is Enabled on the SRX5K-MPC

Purpose Verify that promiscuous mode is enabled on the interface.

Action From operational mode, enter the **show interfaces** command.


```
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
    Bit errors           Seconds
    Errored blocks       0
                          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: Sendbroadcast-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 0x40040000 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
```

Meaning The **Interface flags: Promiscuous** field shows that promiscuous mode is enabled on the interface.

Verifying the Status of Promiscuous Mode

Purpose Verify that promiscuous mode works on the **et-4/0/0** interface.

Action 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	

Meaning 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.

Verifying That Promiscuous Mode Is Disabled

Purpose Verify that disabled promiscuous mode works on the **et-4/0/0** interface.

Action Send traffic into the **et-4/0/0** interface with a MAC address that is different from the interface MAC address and turn off 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	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	

Meaning 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 Promiscuous Mode on Ethernet Interface on page 263](#)
- [Enabling and Disabling Promiscuous Mode on Ethernet Interfaces \(CLI Procedure\) on page 270](#)

Configuring Port Mirroring on SRX Devices

To configure port mirroring on an SRX device, you must first configure the **forwarding-options** and **interfaces** at the **[edit]** hierarchy level.

You must configure the **forwarding-options** statement to define an instance of the **mirror-to** port for port mirroring and also configure the interface to be mirrored.



NOTE: The mirrored port and the mirror-to port must be under the same Broadcom chipset in a I/O card.

To configure port mirroring:

1. Specify the **rate** and **run-length** at the **[edit forwarding-options port-mirroring input]** hierarchy level:



NOTE:

- **rate:** Ratio of packets to be sampled (1 out of *N*) (1 through 65535)
- **run-length:** Number of samples after initial trigger (0 through 20)

```
[edit]
forwarding-options
  port-mirroring {
    input {
      rate number;
      run-length number;
    }
  }
```

2. To send the copies of the packet to the **mirror-to** port, include the **interface *intf-name*** statement at the **[edit forwarding-options port-mirroring family any output]** hierarchy level.

```
output {
```

```
interface intf-name;
}
```



NOTE: Port mirroring on SRX devices uses family any to transfer the mirror-to port information to the Packet Forwarding Engine (PFE). The mirroring engine copies all the packets from mirrored port to the mirror-to port.



NOTE: You can configure an instance clause to specify multiple mirror-to ports.

To mirror an interface, include the port-mirror-instance statement at the [edit interface mirrored-intf-name] hierarchy level.

The mirrored interface is configured with an instance name, defined in the forwarding-options. The mirrored port and the mirror-to port are linked through that instance.

```
instance {
  inst-name {
    input {
      rate number;
      run-length number;
    }
    family any {
      output {
        interface intf-name;
      }
    }
  }
}
interfaces
  mirrored-intf-name {
    port-mirror-instance instance-name;
  }
```



NOTE: Port mirroring on SRX devices does not differentiate the traffic direction, but mirrors the ingress and egress samples together.

A sample configuration for port mirroring is shown below:

```
mirror port ge-1/0/2 to port ge-1/0/9.0
forwarding-options
  port-mirroring {
    input {
      rate 1;
      run-length 10;
    }
  }
```

```
family any {  
  output {  
    interface ge-1/0/9.0;  
  }  
}  
instance {  
  inst1 {  
    input {  
      rate 1;  
      run-length 10;  
    }  
    family any {  
      output {  
        interface ge-1/0/9.0;  
      }  
    }  
  }  
}  
interfaces {  
  ge-1/0/2 {  
    port-mirror-instance inst1;  
  }  
}
```

Related Documentation [Understanding Port Mirroring on SRX Devices on page 264](#)

Configuring Aggregated Ethernet Interfaces

- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
- [Understanding the Aggregated Ethernet Interfaces Device Count on page 282](#)
- [Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device on page 283](#)
- [Understanding Physical Interfaces for Aggregated Ethernet Interfaces on page 284](#)
- [Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces on page 285](#)
- [Understanding Aggregated Ethernet Interface Link Speed on page 286](#)
- [Example: Configuring Aggregated Ethernet Link Speed on page 286](#)
- [Understanding Minimum Links for Aggregated Ethernet Interfaces on page 287](#)
- [Example: Configuring Aggregated Ethernet Minimum Links on page 288](#)
- [Understanding Aggregated Ethernet Interface Removal on page 289](#)
- [Example: Deleting Aggregated Ethernet Interfaces on page 289](#)
- [Example: Deleting Aggregated Ethernet Interface Contents on page 290](#)
- [Verifying Aggregated Ethernet Interfaces on page 291](#)
- [Understanding VLAN Tagging for Aggregated Ethernet Interfaces on page 293](#)
- [Understanding Promiscuous Mode for Aggregated Ethernet Interfaces on page 293](#)

Understanding Aggregated Ethernet Interfaces

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 Feature Guide*.

This topic contains the following sections:

- [LAGs on page 280](#)
- [LACP on page 280](#)

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 Feature 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.

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.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
- [Understanding LACP on Standalone Devices on page 295](#)
- [Understanding LACP on Chassis Clusters on page 301](#)
- [Understanding VLAN Tagging for Aggregated Ethernet Interfaces on page 293](#)

- [Understanding Promiscuous Mode for Aggregated Ethernet Interfaces on page 293](#)

Aggregated Ethernet Interfaces Configuration Overview



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” on page 283](#).
2. Associate a physical interface with the aggregated Ethernet interface. See [“Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces” on page 285](#).
3. (Optional) Set the required link speed for all the interfaces included in the bundle. See [“Example: Configuring Aggregated Ethernet Link Speed” on page 286](#).
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 288](#).
5. (Optional) Enable or disable VLAN tagging. See [“Understanding VLAN Tagging for Aggregated Ethernet Interfaces” on page 293](#).
6. (Optional) Enable promiscuous mode. See [“Understanding Promiscuous Mode for Aggregated Ethernet Interfaces” on page 293](#).

Related Documentation

- [Ethernet Switching Feature Guide](#)
- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Example: Configuring Link Aggregation Control Protocol \(CLI Procedure\) on page 296](#)
- [Example: Configuring LACP on Chassis Clusters on page 304](#)

Understanding the Aggregated Ethernet Interfaces Device Count

By default, no aggregated Ethernet interfaces are created. You must set the number of aggregated Ethernet interfaces on the routing device before you can configure them. Once you set the device count, the system creates that number of empty aggregated Ethernet interfaces. A globally unique MAC address is assigned to every aggregated Ethernet interface. More aggregated Ethernet interfaces can be created by increasing the parameter.

The maximum number of aggregated devices you can configure is 128. The aggregated interfaces are numbered from ae0 through ae127.

Similarly, you can permanently remove an aggregated Ethernet interface from the device configuration by deleting it from the device count. When you reduce the device count, only the aggregated Ethernet interface objects at the end of the list are removed, leaving the newly specified number of interfaces. That is, if you set the device count to 10 and then reduce it to 6, the system removes the last 4 interface objects from the list.



WARNING: Be aware that this approach deletes the aggregated Ethernet interface and *all* of its objects from the device configuration.

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device on page 283](#)
- [Example: Deleting Aggregated Ethernet Interfaces on page 289](#)

Example: Configuring the Number of Aggregated Ethernet Interfaces on a Device

This example shows how to configure the number of aggregated Ethernet interfaces on a device.

- [Requirements on page 283](#)
- [Overview on page 283](#)
- [Configuration on page 283](#)
- [Verification on page 284](#)

Requirements

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

Overview

In this example, you create two aggregate Ethernet interfaces, thereby enabling all the interfaces that you need for your configuration in one step.

Configuration

Step-by-Step Procedure

To configure the number of aggregated Ethernet interfaces on a device:

1. Set the number of aggregated Ethernet interfaces.

[edit]

```
user@host# set chassis aggregated-devices ethernet device-count 2
```

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.

Related Documentation

- [Understanding the Aggregated Ethernet Interfaces Device Count on page 282](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
- [Example: Deleting Aggregated Ethernet Interfaces on page 289](#)
- [Verifying Aggregated Ethernet Interfaces on page 291](#)

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 *aex* 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, 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.
-

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces on page 285](#)

Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces

This example shows how to associate physical interfaces with aggregated Ethernet interfaces.

- [Requirements on page 285](#)
- [Overview on page 285](#)
- [Configuration on page 285](#)
- [Verification on page 285](#)

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” on page 283](#).

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

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.

- Related Documentation**
- [Understanding Physical Interfaces for Aggregated Ethernet Interfaces on page 284](#)
 - [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
 - [Verifying Aggregated Ethernet Interfaces on page 291](#)

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.

- Related Documentation**
- [Understanding Aggregated Ethernet Interfaces on page 279](#)
 - [Example: Configuring Aggregated Ethernet Link Speed on page 286](#)
 - [Understanding Minimum Links for Aggregated Ethernet Interfaces on page 287](#)

Example: Configuring Aggregated Ethernet Link Speed

This example shows how to configure the aggregated Ethernet link speed.

- [Requirements on page 286](#)
- [Overview on page 287](#)
- [Configuration on page 287](#)
- [Verification on page 287](#)

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](#)” on page 283.
- Associate physical interfaces with the aggregated Ethernet Interfaces. See “[Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces](#)” on page 285.

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

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.

Related Documentation

- [Understanding Aggregated Ethernet Interface Link Speed on page 286](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
- [Verifying Aggregated Ethernet Interfaces on page 291](#)

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.

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Example: Configuring Aggregated Ethernet Minimum Links on page 288](#)
- [Understanding Aggregated Ethernet Interface Link Speed on page 286](#)

Example: Configuring Aggregated Ethernet Minimum Links

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 on page 288](#)
- [Overview on page 288](#)
- [Configuration on page 288](#)
- [Verification on page 288](#)

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” on page 283](#).
- Associate physical interfaces with the aggregated Ethernet Interfaces. See [“Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces” on page 285](#).
- Configure the aggregated Ethernet link speed. See [“Example: Configuring Aggregated Ethernet Link Speed” on page 286](#).

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

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.

Related Documentation

- [Understanding Aggregated Ethernet Interface Link Speed on page 286](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
- [Verifying Aggregated Ethernet Interfaces on page 291](#)

Understanding Aggregated Ethernet Interface Removal

You can delete an aggregated Ethernet interface from the interface configuration. Junos OS removes the configuration statements related to **aex** and sets this interface to the down state. The deleted aggregated Ethernet interface still exists, but it becomes an empty interface.

- Related Documentation**
- [Understanding Aggregated Ethernet Interfaces on page 279](#)
 - [Example: Deleting Aggregated Ethernet Interfaces on page 289](#)
 - [Example: Deleting Aggregated Ethernet Interface Contents on page 290](#)

Example: Deleting Aggregated Ethernet Interfaces

This example shows how to delete aggregated Ethernet interfaces using the device count.

- [Requirements on page 289](#)
- [Overview on page 289](#)
- [Configuration on page 289](#)
- [Verification on page 290](#)

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” on page 283](#).

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

- 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.

Related Documentation

- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
- [Example: Deleting Aggregated Ethernet Interface Contents on page 290](#)
- [Verifying Aggregated Ethernet Interfaces on page 291](#)

Example: Deleting Aggregated Ethernet Interface Contents

This example shows how to delete the contents of an aggregated Ethernet interface.

- [Requirements on page 290](#)
- [Overview on page 290](#)
- [Configuration on page 290](#)
- [Verification on page 291](#)

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” on page 283](#).
- Associate a physical interface with the aggregated Ethernet interface. See [“Example: Associating Physical Interfaces with Aggregated Ethernet Interfaces” on page 285](#).
- Set the required link speed for all the interfaces included in the bundle. See [“Example: Configuring Aggregated Ethernet Link Speed” on page 286](#).
- 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 288](#).

Overview

In this example, you delete the contents of the ae4 aggregated Ethernet interface, which sets it to the down state.

Configuration

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.

- Related Documentation**
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
 - [Example: Deleting Aggregated Ethernet Interfaces on page 289](#)
 - [Verifying Aggregated Ethernet Interfaces on page 291](#)

Verifying Aggregated Ethernet Interfaces

- [Verifying Aggregated Ethernet Interfaces \(terse\) on page 291](#)
- [Verifying Aggregated Ethernet Interfaces \(extensive\) on page 292](#)

Verifying Aggregated Ethernet Interfaces (terse)

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.

- See Also**
- [Verifying Aggregated Ethernet Interfaces \(extensive\) on page 292](#)
 - [Understanding Aggregated Ethernet Interfaces on page 279](#)
 - [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)

Verifying Aggregated Ethernet Interfaces (extensive)

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.

See Also

- [Verifying Aggregated Ethernet Interfaces \(terse\) on page 291](#)
- [Understanding Aggregated Ethernet Interfaces on page 279](#)

- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)

**Related
Documentation**

- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)

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**

**Related
Documentation**

- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)

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.

**Related
Documentation**

- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)

Configuring Link Aggregation Control Protocol

- [Understanding LACP on Standalone Devices on page 295](#)
- [Example: Configuring Link Aggregation Control Protocol \(CLI Procedure\) on page 296](#)
- [Verifying LACP on Standalone Devices on page 300](#)
- [Understanding LACP on Chassis Clusters on page 301](#)
- [Example: Configuring LACP on Chassis Clusters on page 304](#)
- [Verifying LACP on Redundant Ethernet Interfaces on page 306](#)
- [LAG and LACP Support on SRX5000 Line Devices with I/O Cards \(IOCs\) on page 307](#)
- [Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3 on page 309](#)
- [Example: Configuring Aggregated Ethernet Device with LAG and LACP on a Security Device \(CLI Procedure\) on page 313](#)

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: On SRX5400, SRX5600, and SRX5800, the LACP is not supported on Layer 2 interfaces.

Related Documentation

- [Understanding Aggregated Ethernet Interfaces on page 279](#)
- [Understanding LACP on Chassis Clusters on page 301](#)
- [Example: Configuring Link Aggregation Control Protocol \(CLI Procedure\) on page 296](#)

Example: Configuring Link Aggregation Control Protocol (CLI Procedure)

This example shows how to configure LACP.

- [Requirements on page 296](#)
- [Overview on page 296](#)
- [Configuration on page 296](#)
- [Verification on page 298](#)

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

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

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.

Related Documentation

- *Understanding Link Aggregation Control Protocol*
- *Ethernet Ports Switching Overview for Security Devices*

Verifying LACP on Standalone Devices

- [Verifying LACP Statistics on page 300](#)
- [Verifying LACP Aggregated Ethernet Interfaces on page 300](#)

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-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
```

- See Also**
- [Understanding LACP on Standalone Devices on page 295](#)
 - [Example: Configuring Link Aggregation Control Protocol \(CLI Procedure\) on page 296](#)

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-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 if it is collecting and distributing packets.

See Also • [Verifying LACP on Redundant Ethernet Interfaces on page 306](#)

Related Documentation • [Example: Configuring Link Aggregation Control Protocol \(CLI Procedure\) on page 296](#)
 • [Verifying LACP on Redundant Ethernet Interfaces on page 306](#)

Understanding LACP on Chassis Clusters

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.

LAGs can be established across nodes in a chassis cluster to provide increased interface bandwidth and link availability.

The Link Aggregation Control Protocol (LACP) provides additional functionality for LAGs. 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.

You configure LACP on a redundant Ethernet interface by setting the LACP mode for the parent link with the **lACP** statement. The LACP mode can be off (the default), active, or passive.

This topic contains the following sections:

- [Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups on page 302](#)
- [Sub-LAGs on page 302](#)
- [Supporting Hitless Failover on page 303](#)
- [Managing Link Aggregation Control PDUs on page 303](#)

Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups

A redundant Ethernet interface has active and standby links located on two nodes in a chassis cluster. All active links are located on one node, and all standby links are located on the other node. You can configure up to eight active links and eight standby links per node.

When at least two physical child interface links from each node are included in a redundant Ethernet interface configuration, the interfaces are combined within the redundant Ethernet interface to form a redundant Ethernet interface LAG.

Having multiple active redundant Ethernet interface links reduces the possibility of failover. For example, when an active link is out of service, all traffic on this link is distributed to other active redundant Ethernet interface links, instead of triggering a redundant Ethernet active/standby failover.

Aggregated Ethernet interfaces, known as local LAGs, are also supported on either node of a chassis cluster but cannot be added to redundant Ethernet interfaces. Likewise, any child interface of an existing local LAG cannot be added to a redundant Ethernet interface, and vice versa. The total maximum number of combined individual node LAG interfaces (ae) and redundant Ethernet (reth) interfaces per cluster is 128.

However, aggregated Ethernet interfaces and redundant Ethernet interfaces can coexist, because the functionality of a redundant Ethernet interface relies on the Junos OS aggregated Ethernet framework.

For more information, see *Understanding Chassis Cluster Redundant Ethernet Interface Link Aggregation Groups*.

Minimum Links

Redundant Ethernet interface configuration includes a **minimum-links** setting that allows you to set a minimum number of physical child links in a redundant Ethernet interface LAG that must be working on the primary node for the interface to be up. The default **minimum-links** value is 1. When the number of physical links on the primary node in a redundant Ethernet interface falls below the **minimum-links** value, the interface might be down even if some links are still working. For more information, see *Example: Configuring Chassis Cluster Minimum Links*.

Sub-LAGs

LACP maintains a point-to-point LAG. Any port connected to the third point is denied. However, a redundant Ethernet interface does connect to two different systems or two remote aggregated Ethernet interfaces by design.

To support LACP on redundant Ethernet interface active and standby links, a redundant Ethernet interface is created automatically to consist of two distinct sub-LAGs, where all active links form an active sub-LAG and all standby links form a standby sub-LAG.

In this model, LACP selection logic is applied and limited to one sub-LAG at a time. In this way, two redundant Ethernet interface sub-LAGs are maintained simultaneously while all the LACP advantages are preserved for each sub-LAG.

It is necessary for the switches used to connect the nodes in the cluster to have a LAG link configured and 802.3ad enabled for each LAG on both nodes so that the aggregate links are recognized as such and correctly pass traffic.



NOTE: The redundant Ethernet interface LAG child links from each node in the chassis cluster must be connected to a different LAG at the peer devices. If a single peer switch is used to terminate the redundant Ethernet interface LAG, two separate LAGs must be used in the switch.

Supporting Hitless Failover

With LACP, the redundant Ethernet interface supports hitless failover between the active and standby links in normal operation. The term *hitless* means that the redundant Ethernet interface state remains up during a failover.

The lacpd process manages both the active and standby links of the redundant Ethernet interfaces. A redundant Ethernet interface state remains up when the number of active up links is equal to or more than the number of minimum links configured. Therefore, to support hitless failover, the LACP state on the redundant Ethernet interface standby links must be collected and distributed before failover occurs.

Managing Link Aggregation Control PDUs

The protocol data units (PDUs) contain information about the state of the link. By default, aggregated and redundant Ethernet links do not exchange link aggregation control PDUs.

You can configure PDUs exchange in the following ways:

- Configure Ethernet links to actively transmit link aggregation control PDUs
- Configure Ethernet links to passively transmit PDUs, sending out link aggregation control PDUs only when they are received from the remote end of the same link

The local end of a child link is known as the actor and the remote end of the link is known as the partner. That is, 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.

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).

For more information, see [“Example: Configuring LACP on Chassis Clusters” on page 304](#).

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.

Example: Configuring LACP on Chassis Clusters

This example shows how to configure LACP on chassis clusters.

- [Requirements on page 304](#)
- [Overview on page 304](#)
- [Configuration on page 305](#)
- [Verification on page 305](#)

Requirements

Before you begin:

Complete the tasks such as enabling the chassis cluster, configuring interfaces and redundancy groups. See *SRX Series Chassis Cluster Configuration Overview* and *Example: Configuring Chassis Cluster Redundant Ethernet Interfaces* for more details.

Overview


You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle. You configure LACP on a redundant Ethernet interface of SRX series device in chassis cluster.

In this example, you set the LACP mode for the reth1 interface to active and set the link aggregation control PDU transmit interval to slow, which is every 30 seconds.

When you enable LACP, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them (sending out LACP PDUs only when they receive them from another link). One side of the link must be configured as active for the link to be up.

[Figure 19 on page 304](#) shows the topology used in this example.

Figure 19: Topology for LAGs Connecting SRX Series Devices in Chassis Cluster to an EX Series Switch



In the [Figure 19 on page 304](#), the ge-3/0/0 interface on SRX Series device is connected to ge-0/0/0 interface on EX Series switch and the ge-15/0/0 interface is connected to ge-0/0/1 on EX Series switch. For more information on EX Series switch configuration, see [Configuring Aggregated Ethernet LACP \(CLI Procedure\)](#).

Configuration

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 *CLI User Guide*.

To configure LACP on chassis clusters:

1. Bind redundant child physical interfaces to reth1.

```
[edit interfaces]
user@host# set interfaces ge-3/0/0 gigether-options redundant-parent reth1
user@host# set interfaces ge-3/0/1 gigether-options redundant-parent reth1
user@host# set interfaces ge-15/0/0 gigether-options redundant-parent reth1
user@host# set interfaces ge-15/0/1 gigether-options redundant-parent reth1
```
2. Add reth1 to redundancy group 1.

```
[edit interfaces]
user@host# set interfaces reth1 redundant-ether-options redundancy-group 1
```
3. Set the LACP on reth1.

```
[edit interfaces]
user@host# set interfaces reth1 redundant-ether-options lacp active
user@host# set interfaces reth1 redundant-ether-options lacp periodic slow
```
4. Assign an IP address to reth1.

```
[edit interfaces]
user@host# set interfaces reth1 unit 0 family inet address 192.168.2.1/24
```
5. If you are done configuring the device, commit the configuration.

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

Verification

Verifying LACP on Redundant Ethernet Interfaces

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces reth1** command.

```
user@host> show lacp interfaces reth1
Aggregated interface: reth1
  LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
```

```

ge-15/0/0    Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-15/0/0    Partner   No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-15/0/1    Actor    No    No    Yes    Yes    Yes    Yes    Fast    Active
ge-15/0/1    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
LACP protocol: Receive State Transmit State Mux State
ge-15/0/0      Current    Fast periodic Collecting distributing
ge-15/0/1      Current    Fast periodic Collecting distributing
ge-3/0/0       Current    Fast periodic Collecting distributing
ge-3/0/1       Current    Fast periodic Collecting distributing
{primary:node1}

```

The output shows redundant Ethernet interface information, such as the following:

- 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.

Verifying LACP on Redundant Ethernet Interfaces

Purpose Display LACP status information for redundant Ethernet interfaces.

Action From operational mode, enter the **show lacp interfaces reth0** command.

```

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}

```

The output shows redundant Ethernet interface information, such as the following:

- 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.

Related Documentation

- [Example: Configuring LACP on Chassis Clusters on page 304](#)
- [Verifying LACP on Standalone Devices on page 300](#)

LAG and LACP Support on SRX5000 Line Devices with I/O Cards (IOCs)

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 not 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.
-

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.

Related Documentation

- [Aggregated Ethernet Interfaces Configuration Overview on page 282](#)
- [Example: Configuring Link Aggregation Control Protocol \(CLI Procedure\) on page 296](#)
- [Example: Configuring LACP on Chassis Clusters on page 304](#)

Example: Configuring LAG Interface on an SRX5000 Line Device with IOC2 or IOC3

Starting in Junos OS Release 15.1X49-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 on page 309](#)
- [Overview on page 309](#)
- [Configuration on page 310](#)
- [Verification on page 312](#)

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 *Example: Configuring SRX5K-MPC3-100G10G (IOC3) and SRX5K-MPC3-40G10G (IOC3) on an SRX5000 Line Device to Support 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

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
```

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

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	CoI	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  Co1  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, IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface.

Related Documentation

- [Understanding LACP on Chassis Clusters on page 301](#)
- [Verifying LACP on Redundant Ethernet Interfaces on page 306](#)

Example: Configuring Aggregated Ethernet Device with LAG and LACP on a Security Device (CLI Procedure)

- [Requirements on page 313](#)
- [Overview on page 313](#)
- [Configuration on page 314](#)
- [Verification on page 315](#)

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

This example shows the configuration of aggregated Ethernet (ae) devices with LAG and LACP.

Configuration

Step-by-Step Procedure

To configure LAG:

1. Configure the number of aggregated Ethernet interfaces with LAG interface that you need to create. Set the device-count option to 5.

[edit]
user@host# set chassis aggregated-devices ethernet device-count 5
2. Add a port to the aggregated Ethernet interface with LAG.

[edit]
user@host# set interfaces ge-2/0/1 ether-options 802.3ad ae0
user@host# set interfaces ge-2/0/2 ether-options 802.3ad ae0
3. Configure LACP for the aggregated Ethernet interface with LAG.

[edit]
user@host# set interfaces ae0 aggregated-ether-options lacp active
4. Configure family Ethernet switching for the aggregated Ethernet interface with LAG.

[edit]
user@host# set interfaces ae0 unit 0 family ethernet-switching
5. Configure the VLAN vlan20 with VLAN ID 20.

[edit]
user@host# set vlans vlan20 vlan-id 20
6. Add the aggregated Ethernet interface to the VLAN.

[edit]
user@host# set vlans vlan20 interface ae0
7. Check the configuration by entering the **show vlans** and **show interfaces** commands

```
user@host# show vlans
vlan20 {
  vlan-id 20;
  interface {
    ae0.0;
  }
}

user@host# show interfaces
ge-2/0/1 {
  ether-options {
    802.3ad ae0;
  }
}
ge-2/0/2 {
  ether-options {
    802.3ad ae0;
```

```

    }
  }
  ae0 {
    aggregated-ether-options {
      lacp {
        active;
      }
    }
    unit 0 {
      family ethernet-switching;
    }
  }
}

```

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

```

[edit]
user@host# commit

```



NOTE: Likewise, you can configure other devices with LAG and LACP.

Verification

Verifying Aggregated Ethernet Interface with LAG and LACP

Purpose Verify that you can configure aggregated Ethernet interfaces with LAG and LACP.

Action From configuration mode, enter the **show lacp interfaces** to view the LACP interfaces.

```

user@host# run show lacp interfaces
Aggregated interface: ae0
LACP state:

```

	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
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-2/0/2	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
ge-2/0/2	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active

```

LACP protocol:

```

	Receive State	Transmit State	Mux State
ge-2/0/1	Current	Fast periodic	Collecting distributing
ge-2/0/2	Current	Fast periodic	Collecting distributing

From configuration mode, enter the **show vlans** command to view the VLAN interfaces.

```

user@host# run show vlans

```

Name	Tag	Interfaces
default	1	None
vlan20	20	ae0.0

From configuration mode, enter the **show interfaces (interface name)** command to view the status of the ge-2/0/1 and ge-2/0/2 interfaces.

```

user@host# run show interfaces ge-2/0/1 terse

```

Interface	Admin	Link Proto	Local	Remote
ge-2/0/1	up	ethernet	10.0.0.1	10.0.0.2
ge-2/0/2	up	ethernet	10.0.0.3	10.0.0.4

```
ge-2/0/1          up    up
ge-2/0/1.0        up    up    aenet    --> ae0.0

user@host# run show interfaces ge-2/0/2 terse
Interface        Admin Link Proto  Local          Remote
ge-2/0/2          up    up
ge-2/0/2.0        up    up    aenet    --> ae0.0
```

Meaning The output shows the aggregated Ethernet Interface with LAG and LACP is configured.

- Related Documentation**
- [Understanding Aggregated Ethernet Interfaces on page 279](#)
 - [Understanding LACP on Standalone Devices on page 295](#)
 - [Example: Configuring Link Aggregation Control Protocol \(CLI Procedure\) on page 296](#)

Configuring Gigabit Ethernet Physical Interface Modules

- [Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM on page 317](#)
- [Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface on page 319](#)
- [Understanding the 2-Port 10-Gigabit Ethernet XPIM on page 325](#)
- [Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 328](#)
- [Understanding the 8-Port Gigabit Ethernet SFP XPIM on page 332](#)
- [Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs on page 334](#)

Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM

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 on page 317](#)
- [Interface Names and Settings on page 318](#)
- [Available Link Speeds and Modes on page 318](#)
- [Link Settings on page 319](#)

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.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Example: Configuring the 1-Port Gigabit Ethernet SFP Mini-PIM Interface on page 319](#)

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

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

- [Requirements on page 319](#)
- [Overview on page 320](#)
- [Configuration on page 320](#)
- [Verification on page 323](#)

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 265](#).

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

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 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 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 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 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 *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

Confirm that the configuration is working properly.

- [Verifying That the Correct Hardware Is Installed on page 323](#)
- [Verifying the FPC Status on page 324](#)
- [Verifying the Interface Settings on page 325](#)

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
```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)
			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: Sendbroadcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 14.1.1.1/24, 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

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Understanding the 1-Port Gigabit Ethernet SFP Mini-PIM on page 317](#)

Understanding the 2-Port 10-Gigabit Ethernet XPIM

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 on page 326](#)
- [Interface Names and Settings on page 327](#)
- [Copper and Fiber Operating Modes on page 327](#)
- [Link Speeds on page 327](#)
- [Link Settings on page 328](#)

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.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 328](#)

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

This example shows how to perform basic configuration for the 1-Port Gigabit Ethernet SFP Mini-PIM.

- [Requirements on page 328](#)
- [Overview on page 328](#)
- [Configuration on page 329](#)
- [Verification on page 330](#)

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 265.

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

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 *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

Confirm that the configuration is working properly.

- [Verifying That the Correct Hardware Is Installed on page 330](#)
- [Verifying the FPC Status on page 331](#)
- [Verifying the Interface Settings on page 331](#)

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: Sendbroadcast-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

**Related
Documentation**

- [Understanding the 2-Port 10-Gigabit Ethernet XPIM on page 325](#)
- [Understanding Ethernet Interfaces on page 259](#)

Understanding the 8-Port Gigabit Ethernet SFP XPIM

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.

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.

Related Documentation

- [Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs on page 334](#)

Example: Configuring 8-Port Gigabit Ethernet SFP XPIMs

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 on page 335](#)
- [Overview and Topology on page 335](#)
- [Configuration on page 336](#)
- [Verification on page 340](#)

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:

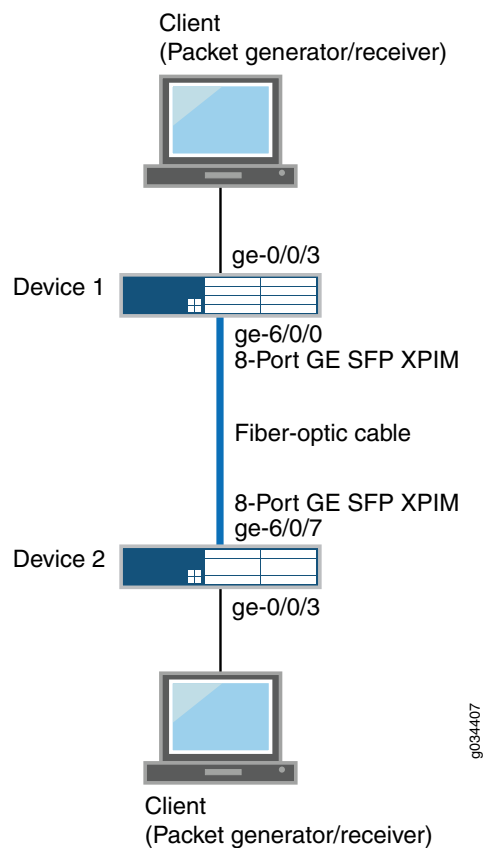
- Establish basic connectivity. See the Getting Started Guide for your device.
- Configure network interfaces as necessary. See [“Example: Creating an Ethernet Interface” on page 265](#).

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 20 on page 336](#) shows the topology used in this example.

Figure 20: Basic Back-to-Back Device Configuration



Configuration

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

```
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 10.1.1.2/24

```



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

Confirm that the configuration is working properly.

- [Verifying the Hardware was Properly Installed on page 341](#)
- [Verifying the FPC Status on page 342](#)
- [Verifying Interface Link Status on Device 1 on page 342](#)
- [Verifying the Interface Settings on Device 1 on page 343](#)

- [Verifying Interface Link Status on Device 2 on page 346](#)
- [Verifying the Interface Settings on Device 2 on page 346](#)

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                               AAK0059
System IO      REV 08    710-023209   AAJ9290        SRXSME System IO
Routing Engine REV 13    750-023223   AAJ1987        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
  PIC 0
FPC 1          REV 03    750-038290   AADL2016       4x GE Base PIC
FPC 5
  PIC 0
FPC 6          REV 03    750-037551   AAEC8065       8x GE SFP gPIM
  PIC 0
    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:
    Total octets          Receive      Transmit
    Total packets        268          268
    Unicast packets       3            3
    Broadcast packets     3            2
    Multicast packets     0            1
    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          low
    0 best-effort             95          950000000      95           0
none
    3 network-control         5           500000000       5           0
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: Sendbroadcast-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.

Related Documentation

- [Understanding the 8-Port Gigabit Ethernet SFP XPIM on page 332](#)

Configuring Port Mirroring

- [Understanding Port Mirroring on SRX Devices on page 351](#)
- [Configuring Port Mirroring on SRX Devices on page 352](#)

Understanding Port Mirroring on SRX Devices

Port mirroring copies packets entering or exiting a port and sends the copies to a local interface for monitoring. Port mirroring is used to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Port mirroring is used to send a copy of all the packets or only the sampled packets seen on a port to a network monitoring connection. You can mirror the packets either on the incoming port (ingress port mirroring) or the outgoing port (egress port mirroring).



NOTE: Port mirroring is supported only on the SRX devices with the following I/O cards:

- SRX1K-SYSIO-GE
- SRX1K-SYSIO-XGE
- SRX3K-SFB-12GE
- SRX3K-2XGE-XFP
- SRX5K-FPC-IOC Flex I/O

On SRX devices, all packets passing through the **mirrored** port are copied and sent to the specified **mirror-to** port. These ports must be on the same Broadcom chipset in the I/O cards.



NOTE: On SRX devices, port mirroring works on physical interfaces only.

Related Documentation

- [Configuring Port Mirroring on SRX Devices on page 275](#)

Configuring Port Mirroring on SRX Devices

To configure port mirroring on an SRX device, you must first configure the **forwarding-options** and **interfaces** at the **[edit]** hierarchy level.

You must configure the **forwarding-options** statement to define an instance of the **mirror-to** port for port mirroring and also configure the interface to be mirrored.



NOTE: The mirrored port and the mirror-to port must be under the same Broadcom chipset in a I/O card.

To configure port mirroring:

1. Specify the **rate** and **run-length** at the **[edit forwarding-options port-mirroring input]** hierarchy level:



NOTE:

- **rate:** Ratio of packets to be sampled (1 out of *N*) (1 through 65535)
- **run-length:** Number of samples after initial trigger (0 through 20)

```
[edit]
  forwarding-options
    port-mirroring {
      input {
        rate number;
        run-length number;
      }
    }
```

2. To send the copies of the packet to the **mirror-to** port, include the **interface *intf-name*** statement at the **[edit forwarding-options port-mirroring family any output]** hierarchy level.

```
  output {
    interface intf-name;
  }
```



NOTE: Port mirroring on SRX devices uses **family any** to transfer the **mirror-to** port information to the Packet Forwarding Engine (PFE). The mirroring engine copies all the packets from mirrored port to the **mirror-to** port.



NOTE: You can configure an instance clause to specify multiple mirror-to ports.

To mirror an interface, include the `port-mirror-instance` statement at the [edit interface mirrored-intf-name] hierarchy level.

The mirrored interface is configured with an instance name, defined in the forwarding-options. The mirrored port and the mirror-to port are linked through that instance.

```
instance {
  inst-name {
    input {
      rate number;
      run-length number;
    }
    family any {
      output {
        interface intf-name;
      }
    }
  }
}
interfaces
  mirrored-intf-name {
    port-mirror-instance instance-name;
  }
```



NOTE: Port mirroring on SRX devices does not differentiate the traffic direction, but mirrors the ingress and egress samples together.

A sample configuration for port mirroring is shown below:

```
mirror port ge-1/0/2 to port ge-1/0/9.0
forwarding-options
  port-mirroring {
    input {
      rate 1;
      run-length 10;
    }
    family any {
      output {
        interface ge-1/0/9.0;
      }
    }
  }
  instance {
    inst1 {
      input {
        rate 1;
        run-length 10;
      }
```

```
        family any {
            output {
                interface ge-1/0/9.0;
            }
        }
    }
}
interfaces {
    ge-1/0/2 {
        port-mirror-instance inst1;
    }
}
```

Related Documentation [Understanding Port Mirroring on SRX Devices on page 264](#)

Configuring Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 355](#)
- [Example: Configuring Ethernet OAM Link Fault Management on a Security Device on page 357](#)
- [Example: Configuring Remote Loopback Mode on VDSL Interfaces on a Security Device on page 361](#)

Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways

Starting in Junos OS Release 15.1X49-D70, Ethernet OAM link fault management for SRX Series services gateways is supported on SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500 devices.

The Ethernet interfaces on SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

The following OAM LFM features are supported:

- **Discovery and link monitoring**—The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The device performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote fault detection—Remote fault detection uses flags and events. Flags convey Link Fault (a loss of signal), Dying Gasp (an unrecoverable condition such as a power failure), and Critical Event (an unspecified vendor-specific critical event). You can specify the periodic OAM PDU sending interval for fault detection. SRX Series devices use the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.
- Remote loopback—Remote loopback mode ensures link quality between the device and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote data terminal equipment (DTE) into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

Table 35 on page 356 lists the interfaces modes supported.

Table 35: Supported Interface Modes

Interfaces	Mode
Physical interface (fe/ge)	Family <ul style="list-style-type: none"> • ccc • ethernet-switching • inet6 • inet • iso • mpls • tcc
	IFD encapsulations <ul style="list-style-type: none"> • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • ethernet-tcc • extended-vlan-tcc

Table 35: Supported Interface Modes (continued)

Interfaces	Mode
Aggregated Ethernet interface (Static or LACP lag)	Family <ul style="list-style-type: none"> • ethernet-switching • inet • mpls • iso • inet6
	IFD encapsulations <ul style="list-style-type: none"> • ethernet-ccc • extended-vlan-ccc (IFD vlan-tagging mode) • vlan-ccc

**Related
Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on a Security Device on page 357](#)

Example: Configuring Ethernet OAM Link Fault Management on a Security Device

Starting in Junos OS Release 15.1X49-D70, configuring Ethernet OAM link fault management is supported on SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500 devices.

The Ethernet interfaces on the SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet or Fast Ethernet interface:

- [Requirements on page 357](#)
- [Overview on page 358](#)
- [Configuration on page 358](#)
- [Verification on page 360](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1 R2 or later for SRX Series Services Gateways
- Any two models of SRX Series devices connected directly

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 265](#).
- Ensure that you configure the interfaces as per the interface modules listed in [“Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways” on page 355](#)


Overview

The Ethernet interfaces on the SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two SRX Series devices connected directly. Before you begin configuring Ethernet OAM LFM on these two devices, connect the two devices directly through supported interfaces. See [“Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways” on page 355](#).

[Figure 21 on page 358](#) shows the topology used in this example.

Figure 21: Ethernet LFM with SRX Series Devices



NOTE: For more information about configuring Ethernet OAM Link Fault Management, see [Junos® OS Ethernet Interfaces](#).

Configuration

To configure Ethernet OAM LFM, perform these tasks:

- [Configuring Ethernet OAM Link Fault Management on Device 1 on page 358](#)
- [Configuring Ethernet OAM Link Fault Management on Device 2 on page 359](#)

Configuring Ethernet OAM Link Fault Management on Device 1

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 protocols oam ethernet link-fault-management interface ge-0/0/0
set protocols oam ethernet link-fault-management interface ge-0/0/0 pdu-interval 800
set protocols oam ethernet link-fault-management interface ge-0/0/0 link-discovery
active

```

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 Ethernet OAM LFM on device 1:

1. Enable IEEE 802.3ah OAM support.

```

[edit protocols oam ethernet link-fault-management]
user@device1# set interface ge-0/0/0

```
2. Set the periodic OAM PDU-sending interval (in milliseconds) for fault detection.

```

[edit protocols oam ethernet link-fault-management]
user@device1# set interface pdu-interval 800

```
3. Specify that the interface initiates the discovery process.

```

[edit protocols oam ethernet link-fault-management]
user@device1# set interface ge-0/0/0 link-discovery active

```

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

```

[edit]
user@device1# show protocols
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/0 {
          pdu-interval 800;
          link-discovery active;
        }
      }
    }
  }
}

```

Configuring Ethernet OAM Link Fault Management on Device 2

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 protocols oam ethernet link-fault-management interface ge-0/0/1
set protocols oam ethernet link-fault-management interface ge-0/0/1 pdu-interval 800
set protocols oam ethernet link-fault-management interface ge-0/0/1 negotiation-options
allow-remote-loopback
```

**Step-by-Step
Procedure**

To configure Ethernet OAM LFM on device 2:

1. Enable OAM on the peer interface.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1
```
2. Set the periodic OAM PDU-sending interval (in milliseconds) for fault detection.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1 pdu-interval 800
```
3. Enable remote loopback support for the local interface.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

Results

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

```
[edit]
user@device2# show protocols
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/1 {
          negotiation-options {
            allow-remote-loopback;
          }
        }
      }
    }
  }
}
```

Verification

Verify the OAM LFM Configuration

Purpose Verify that OAM LFM is configured properly.

Action From operational mode, enter the **show oam ethernet link-fault-management** command.


```
user@device1> show oam ethernet link-fault-management
```

```
Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 2001:bd8:00:31
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
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
```

Meaning The output displays the MAC address and the discovery state is **Send Any** if OAM LFM has been configured properly.

Related Documentation

- [Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 355](#)

Example: Configuring Remote Loopback Mode on VDSL Interfaces on a Security Device

Starting in Junos OS Release 15.1X49-D110, configuring remote loopback mode in Ethernet OAM link fault management (LFM) on a VDSL interface is supported on SRX320, SRX340, SRX345, and SRX550M devices.

This example describes the following configuration scenarios:

Starting in Junos OS Release 12.3X48-D65, configuring remote loopback mode in Ethernet OAM link fault management (LFM) on a VDSL interface is supported on SRX210, SRX220, SRX240, and SRX550 devices.

This example describes the following configuration scenarios:

- Scenario 1: Configuring remote loopback mode on a VDSL interface.
- Scenario 2: Configuring remote loopback mode on a VDSL interface acting as a PPPOE's underlying interface.
- [Requirements on page 361](#)
- [Overview on page 362](#)
- [Configuration for Scenario 1 on page 362](#)
- [Configuration for Scenario 2 on page 363](#)
- [Verification on page 364](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 15.1X49-D110 or later for SRX Series Services Gateways
- An SRX 210/220/240/320/340/345/550/550M device connected with a DSLAM

Before you begin:


- Establish basic connectivity. See the Getting Started Guide for your device.
- Configure network interfaces as necessary. See [“Example: Configuring VDSL2 Interfaces \(Basic\)”](#) on page 220.
- Ensure that you configure the interfaces as per the interface modules listed in [“Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways”](#) on page 355
- Ensure that you configure PPPOE as per the instructions listed in [“Example: Configuring PPPoE Interfaces”](#) on page 393

Overview

This example uses an SRX Series device connected to a DSLAM. Before you begin configuring Ethernet OAM LFM on these two devices, connect the two devices directly through supported interfaces. See [“Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways”](#) on page 355.

Figure 21 on page 358 shows the topology used in this example.

Figure 22: Ethernet LFM with SRX Series Devices



NOTE: For more information about configuring Ethernet OAM Link Fault Management, see [Junos® OS Ethernet Interfaces](#).

Configuration for Scenario 1

To configure remote loopback mode on a VDSL interface, perform these tasks:

- [Configuring Remote Loopback Mode on a VDSL interface of an SRX Series Device](#) on page 362

Configuring Remote Loopback Mode on a VDSL interface of an SRX Series Device

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 protocols oam ethernet link-fault-management interface pt-1/0/0
set protocols oam ethernet link-fault-management interface pt-1/0/0 negotiation-options
allow-remote-loopback
```

- Step-by-Step Procedure** To configure remote loopback mode on a VDSL interface of an SRX Series device:
1. Enable OAM on a VDSL interface.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface pt-1/0/0
```
 2. Enable remote loopback support for the interface.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interface pt-1/0/0 negotiation-options allow-remote-loopback
```

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

```
[edit]
user@device2# show protocols
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface pt-1/0/0 {
          negotiation-options {
            allow-remote-loopback;
          }
        }
      }
    }
  }
}
```

Configuration for Scenario 2

To configure remote loopback mode on a PPPOE's underlying interface, perform these tasks:

- [Configuring Remote Loopback Mode on a PPPOE's underlying interface on page 363](#)

Configuring Remote Loopback Mode on a PPPOE's underlying 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 pp0 unit 0 pppoe-options underlying-interface pt-1/0/0
set protocols oam ethernet link-fault-management interface pt-1/0/0 link-discovery
active
set protocols oam ethernet link-fault-management interface pt-1/0/0 negotiation-options
allow-remote-loopback
```

Step-by-Step Procedure

To configure remote loopback mode on a PPPOE's underlying interface:

1. Create the PPPoE interface pp0 and specify the logical PT interface pt-1/0/0 as the underlying interface.

```
[edit protocols oam ethernet link-fault-management]
user@device2# set interfaces pp0 unit 0 pppoe-options underlying-interface
pt-1/0/0
```

2. Specify that the interface initiates the discovery process.

```
user@device2# set protocols oam ethernet link-fault-management interface
pt-1/0/0 link-discovery active
```

3. Enable remote loopback mode.

```
user@device2# set protocols oam ethernet link-fault-management interface
pt-1/0/0 negotiation-options allow-remote-loopback
```

Results

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

```
[edit]
user@device2# show protocols
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface pt-1/0/0 {
          link-discovery active;
          negotiation-options {
            allow-remote-loopback;
          }
        }
      }
    }
  }
}
```

Verification

Verify the OAM LFM Configuration

Purpose Verify that OAM LFM is configured properly.

Action From operational mode, enter the **show oam ethernet link-fault-management** command.

```
user@device1> show oam ethernet link-fault-management
```

```
Interface: pt-1/0/0.0
Status: Running, Discovery state: Send Any
```

```
Transmit interval: 300ms, PDU threshold: 3 frames, Hold time: 900ms
Peer address: 2001:db8:e5:b9:c8:ed
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Loopback tracking: Disabled, Loop status: Unknown
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: unsupported, Link events: supported
Variable requests: unsupported
```

Meaning The output displays the MAC address and the discovery state is **Send Any** if OAM LFM has been configured properly.

Related Documentation

- [Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 355](#)

Configuring Power over Ethernet

- [Understanding Power over Ethernet on page 367](#)
- [Example: Configuring PoE on All Interfaces on page 370](#)
- [Example: Configuring PoE on an Individual Interface on page 373](#)
- [Example: Disabling a PoE Interface on page 376](#)

Understanding Power over Ethernet

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 SRX Series devices support PoE on Ethernet ports. PoE ports transfer electrical power and data to remote devices over standard twisted-pair cable in an Ethernet network. PoE ports allow you to plug in devices that require both network connectivity and electrical power, such as VoIP and IP phones and wireless LAN access points.

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.

This topic contains the following sections:

- [SRX Series Services Gateway PoE Specifications on page 367](#)
- [PoE Classes and Power Ratings on page 369](#)
- [PoE Options on page 370](#)

SRX Series Services Gateway PoE Specifications

[Table 36 on page 368](#) lists the PoE specifications for the SRX210, SRX240, SRX320, SRX340, and SRX650 devices. (Platform support depends on the Junos OS release in your installation.)

Table 36: PoE Specifications for the SRX210, SRX240, SRX320, SRX340, and SRX650 Devices

Specifications	For SRX210 Device	For SRX240 Device	For SRX320 Device	For SRX340 Device	For SRX650 Device
Supported standards	<ul style="list-style-type: none"> IEEE 802.3 AF Legacy (pre-standards) 	<ul style="list-style-type: none"> IEEE 802.3 AF IEEE 802.3 AT (PoE+) Legacy (pre-standards) 	<ul style="list-style-type: none"> IEEE 802.3 AF Legacy (pre-standards) 	<ul style="list-style-type: none"> IEEE 802.3 AF IEEE 802.3 AT (PoE+) Legacy (pre-standards) 	<ul style="list-style-type: none"> IEEE 802.3 AF IEEE 802.3 AT (PoE+) Legacy (pre-standards)
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 16 Gigabit Ethernet ports (ge-0/0/0 to ge-0/0/15).	Supported on two Gigabit Ethernet ports and two Fast Ethernet ports (ge-0/0/0 to ge-0/0/5).	Supported on all 16 Gigabit Ethernet ports (ge-0/0/0 to ge-0/0/15).	Supported on the following ports: <ul style="list-style-type: none"> Slot 2 or 6 on 16 Gigabit Ethernet ports <ul style="list-style-type: none"> 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 <ul style="list-style-type: none"> ge-2/0/0 to ge-2/0/23 ge-6/0/0 to ge-6/0/23
Total PoE power sourcing capacity	50 W	150 W	50 W	150 W	The 645 watts AC and 645 watts DC power supplies support the following capacities: <ul style="list-style-type: none"> 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.
Default per port power limit	15.4 W	15.4 W	15.4 W	15.4 W	15.4 W
Maximum per port power limit	30 W	30 W	30 W	30 W	30 W

Table 36: PoE Specifications for the SRX210, SRX240, SRX320, SRX340, and SRX650 Devices (continued)

Specifications	For SRX210 Device	For SRX240 Device	For SRX320 Device	For SRX340 Device	For SRX650 Device
Power management modes	<ul style="list-style-type: none"> Static: Power allocated for each interface can be configured. Class: Power allocated for interfaces is based on the class of powered device connected. 	<ul style="list-style-type: none"> Static: Power allocated for each interface can be configured. Class: Power allocated for interfaces is based on the class of powered device connected. 	<ul style="list-style-type: none"> Static: Power allocated for each interface can be configured. Class: Power allocated for interfaces is based on the class of powered device connected. 	<ul style="list-style-type: none"> Static: Power allocated for each interface can be configured. Class: Power allocated for interfaces is based on the class of powered device connected. 	<ul style="list-style-type: none"> 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

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

[Table 37 on page 369](#) lists the classes and their power ratings as specified by the IEEE standards.

Table 37: 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 configuring PoE, you must enable the PoE interface in order 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 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. (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. The default wattage per port is 15.4 watts.
- PoE power consumption logging—Allows logging of per-port PoE power consumption. The telemetry section must be explicitly specified to enable logging. If left unspecified, telemetry is disabled by default. The default telemetry duration is 1 hour. The default telemetry interval is 5 minutes.
- PoE power management mode—Has two modes:
 - 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 standards.
 - 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.
- Reserve power—Reserves the specified amount of power for the gateway in case of a spike in PoE consumption. The default is 0.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)
- [Example: Configuring PoE on All Interfaces on page 370](#)
- [Example: Configuring PoE on an Individual Interface on page 373](#)
- [Example: Disabling a PoE Interface on page 376](#)

Example: Configuring PoE on All Interfaces

This example shows how to configure PoE on all interfaces.

- [Requirements on page 371](#)
- [Overview on page 371](#)
- [Configuration on page 371](#)
- [Verification on page 372](#)

Requirements

Before you begin, configure Ethernet interfaces. See [“Example: Creating an Ethernet Interface” on page 265](#).

Overview

This example shows how to configure PoE on all interfaces on a device. In this example, you set the power port priority to low and the maximum power available to a port to 15.4 watts. Then you enable the PoE power consumption logging with the default telemetry settings, and you set the PoE management mode to static. Finally, you set the reserved power consumption to 15 watts in case of a spike in PoE consumption.

Configuration

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 poe interface all priority low maximum-power 15.4 telemetry
set poe management static guard-band 15
```

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 PoE on all interfaces:

1. Enable PoE.

```
[edit]
user@host# edit poe interface all
```
2. Set the power port priority.

```
[edit poe interface all]
user@host# set priority low
```
3. Set the maximum PoE wattage available for a port.

```
[edit poe interface all]
user@host# set maximum-power 15.4
```
4. Enable logging of PoE power consumption.

```
[edit poe interface all]
user@host# set telemetry
```
5. Set the PoE management mode.

```
[edit]
user@host# set poe management static
```

6. Reserve power wattage in case of a spike in PoE consumption.

[edit]

```
user@host# set poe guard-band 15
```

Results From configuration mode, confirm your configuration by entering the **show poe interface all** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

[edit]

```
user@host# show poe interface all
priority low;
maximum-power 15.4;
telemetries;
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Status of PoE Interfaces

Purpose 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.)

Action From operational mode, enter the **show poe interface all** command.

```
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

The **show poe interface all** command lists PoE interfaces configured on the SRX 240 device, including information on status, priority, power consumption, and class. 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.

Related Documentation

- [Understanding Power over Ethernet on page 367](#)
- [Example: Configuring PoE on an Individual Interface on page 373](#)
- [Example: Disabling a PoE Interface on page 376](#)

Example: Configuring PoE on an Individual Interface

This example shows how to configure PoE on an individual interface.

- [Requirements on page 373](#)
- [Overview on page 373](#)
- [Configuration on page 373](#)
- [Verification on page 374](#)

Requirements

Before you begin:

- Configure Ethernet interfaces. See “[Example: Creating an Ethernet Interface](#)” on page 265.
- Configure PoE on all interfaces. See “[Example: Configuring PoE on All Interfaces](#)” on page 370.

Overview

This example shows how to configure PoE on the ge-0/0/0 interface. In this example, you set the power port priority to high and the maximum power available to a port to 15.4 watts. Then you enable the PoE power consumption logging with the default telemetry settings, and you set the PoE management mode to static. Finally, you set the reserved power to 15 watts in case of a spike in PoE consumption.

Configuration

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 poe interface ge-0/0/0 priority high maximum-power 15.4 telemetry
set poe management static guard-band 15
```

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 PoE:

1. Enable PoE.

```
[edit]
user@host# edit poe interface ge-0/0/0
```
2. Set the power port priority.

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

3. Set the maximum PoE wattage available for a port.

```
[edit poe interface ge-0/0/0]  
user@host# set maximum power 15.4
```
4. Enable logging of PoE power consumption.

```
[edit poe interface ge-0/0/0]  
user@host# set telemetries
```
5. Set the PoE management mode.

```
[edit]  
user@host# set poe management static
```
6. Reserve power wattage in case of a spike in PoE consumption.

```
[edit]  
user@host# set poe guard-band 15
```

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

```
[edit]  
user@host# show poe interface ge-0/0/0  
priority high;  
maximum-power 15.4;  
telemetries;
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the Status of PoE Interfaces on page 374](#)
- [Verifying the Telemetry Data \(History\) for the Specified Interface on page 375](#)
- [Verifying PoE Global Parameters on page 375](#)

Verifying the Status of PoE Interfaces

Purpose Verify that the PoE interfaces on the device are enabled and set to the desired priority settings. (The device used in this example is the SRX240 or SRX340 Services Gateway, depending on the Junos OS release in the installation.)

Action From operational mode, enter the **show poe interface ge-0/0/1** command.

```
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

```

The **show poe interface ge-0/0/1** command lists PoE interfaces configured on the SRX340 device, with their status, priority, power consumption, and class.

Verifying the Telemetry Data (History) for the Specified Interface

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

Action From operational mode, enter the **show poe telemetries interface** command.

For all records:

```

user@host> show poe telemetries interface ge-0/0/1 all
SI 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
SI 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.

Verifying PoE Global Parameters

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

Action From operational mode, enter the **show poe controller** command.

```

user@host> show poe controller
Controller Maximum Power Guard band Management
index power consumption
0 150.0 W 0.0 W 0 W Static

```

The **show poe controller** command lists the global parameters configured on the SRX Series device such as controller index, maximum power, power consumption, guard band, and management mode along with their status.

- Related Documentation**
- [Understanding Power over Ethernet on page 367](#)
 - [Example: Configuring PoE on All Interfaces on page 370](#)
 - [Example: Disabling a PoE Interface on page 376](#)

Example: Disabling a PoE Interface

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

- [Requirements on page 376](#)
- [Overview on page 376](#)
- [Configuration on page 376](#)
- [Verification on page 377](#)

Requirements

Before you begin:

- Configure PoE on all interfaces. See “[Example: Configuring PoE on All Interfaces](#)” on [page 370](#).
- Configure PoE on an individual interface. See “[Example: Configuring PoE on an Individual Interface](#)” on [page 373](#).

Overview

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

Configuration

Step-by-Step Procedure

To disable PoE on interfaces:

1. Disable PoE on all interfaces.

[edit]
user@host# **set poe interface all disable**
2. Disable PoE on a specific interface.

[edit]
user@host# **set poe interface ge-0/0/0 disable**
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 poe interface** command.

Related Documentation

- [Understanding Power over Ethernet on page 367](#)

PART 5

Configuring Interface Encapsulation

- [Interface Encapsulation Overview on page 381](#)
- [Configuring Point-to-Point Protocol over Ethernet on page 389](#)
- [Configuring PPPoE-Based Radio-to-Router Protocol on page 415](#)
- [Configuring R2CP Radio-to-Router Protocol on page 423](#)

Interface Encapsulation Overview

- [Understanding Physical Encapsulation on an Interface on page 381](#)
- [Understanding Frame Relay Encapsulation on an Interface on page 382](#)
- [Understanding Point-to-Point Protocol on page 383](#)
- [Understanding High-Level Data Link Control on page 386](#)

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 382](#).
- Point-to-Point Protocol. See [“Understanding Point-to-Point Protocol” on page 383](#).
- Point-to-Point Protocol over Ethernet. See [“Understanding Point-to-Point Protocol over Ethernet” on page 389](#).
- High-Level Data Link Control. See [“Understanding High-Level Data Link Control” on page 386](#).


Related Documentation

- [Understanding Interfaces on page 3](#)
- [Understanding Frame Relay Encapsulation on an Interface on page 382](#)
- [Understanding Point-to-Point Protocol on page 383](#)
- [Understanding High-Level Data Link Control on page 386](#)

Understanding Frame Relay Encapsulation on an Interface

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 23 on page 382](#) shows a typical Frame Relay network.

Figure 23: Frame Relay Network



[Figure 23 on page 382](#) 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 on page 382](#)
- [Switched and Permanent Virtual Circuits on page 382](#)
- [Data-Link Connection Identifiers on page 383](#)
- [Congestion Control and Discard Eligibility on page 383](#)

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.

Related Documentation

- [Understanding Physical Encapsulation on an Interface on page 381](#)

Understanding Point-to-Point Protocol

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 on page 384](#)
- [PPP Authentication on page 384](#)
- [Network Control Protocols on page 385](#)
- [Magic Numbers on page 385](#)
- [CSU/DSU Devices on page 386](#)

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.

Related Documentation

- [Understanding Physical Encapsulation on an Interface on page 381](#)

Understanding High-Level Data Link Control

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 on page 386](#)
- [HDLC Operational Modes on page 387](#)

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.

Related Documentation

- [Understanding Physical Encapsulation on an Interface on page 381](#)

Configuring Point-to-Point Protocol over Ethernet

- [Understanding Point-to-Point Protocol over Ethernet on page 389](#)
- [Understanding PPPoE Interfaces on page 392](#)
- [Example: Configuring PPPoE Interfaces on page 393](#)
- [Understanding PPPoE Ethernet Interfaces on page 399](#)
- [Example: Configuring PPPoE Encapsulation on an Ethernet Interface on page 400](#)
- [Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 401](#)
- [Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 402](#)
- [Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 404](#)
- [Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 404](#)
- [Understanding CHAP Authentication on a PPPoE Interface on page 407](#)
- [Example: Configuring CHAP Authentication on a PPPoE Interface on page 407](#)
- [Verifying Credit-Flow Control on page 409](#)
- [Verifying PPPoE Interfaces on page 410](#)
- [Verifying R2CP Interfaces on page 410](#)
- [Displaying Statistics for PPPoE on page 411](#)
- [Setting Tracing Options for PPPoE on page 412](#)

Understanding Point-to-Point Protocol over Ethernet

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 on page 390](#)
- [PPPoE Session Stage on page 391](#)

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 secreta

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).

**Related
Documentation**

- [Understanding Physical Encapsulation on an Interface on page 381](#)
- [Understanding PPPoE Interfaces on page 392](#)
- [Understanding PPPoE Ethernet Interfaces on page 399](#)
- [Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 401](#)
- [Understanding CHAP Authentication on a PPPoE Interface on page 407](#)
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 416](#)

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 PPPoE session in the event of a failover.

- Related Documentation**
- [Understanding Point-to-Point Protocol on page 383](#)
 - [Example: Configuring PPPoE Interfaces on page 393](#)

Example: Configuring PPPoE Interfaces

This example shows how to configure a PPPoE interface.

- [Requirements on page 393](#)
- [Overview on page 393](#)
- [Configuration on page 393](#)
- [Disabling the End-of-List Tag on page 397](#)

Requirements

Before you begin, configure an Ethernet interface. See “[Example: Creating an Ethernet Interface](#)” on page 265.

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

- 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
```

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

Confirm that the configuration is working properly.

- [Verifying PPPoE Interfaces on page 395](#)
- [Verifying PPPoE Sessions on page 396](#)
- [Verifying the PPPoE Version on page 396](#)
- [Verifying PPPoE Statistics on page 397](#)

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 interface
- For packet type, the number of packets of each type sent and received during the PPPoE session

Disabling the End-of-List Tag

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.

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.

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.

- Related Documentation
- [Understanding PPPoE Interfaces on page 392](#)

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 24 on page 400](#) shows a typical PPPoE session between a device and an access concentrator on the Ethernet loop.

Figure 24: PPPoE Session on the Ethernet Loop

To configure PPPoE on an Ethernet interface, you configure encapsulation on the logical interface.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 389](#)
- [Example: Configuring PPPoE Encapsulation on an Ethernet Interface on page 400](#)

Example: Configuring PPPoE Encapsulation on an Ethernet Interface

This example shows how to configure PPPoE encapsulation on an Ethernet interface.

- [Requirements on page 400](#)
- [Overview on page 400](#)
- [Configuration on page 400](#)
- [Verification on page 401](#)

Requirements

Before you begin:

- Configure an Ethernet interface. See [“Example: Creating an Ethernet Interface” on page 265](#).
- Configure a PPPoE encapsulation interface. See [“Example: Configuring PPPoE Interfaces” on page 393](#).

Overview

In this example, you configure PPPoE encapsulation on the ge-0/0/1 interface.

Configuration

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.


Related Documentation

- [Understanding PPPoE Ethernet Interfaces on page 399](#)

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 25 on page 401](#) shows a typical PPPoE over ATM session between a device and an access concentrator on an ADSL loop.

Figure 25: 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.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 389](#)
- [Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 402](#)

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

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 on page 402](#)
- [Overview on page 402](#)
- [Configuration on page 402](#)
- [Verification on page 403](#)

Requirements

Before you begin:

- Configure network interfaces. See “[Example: Creating an Ethernet Interface](#)” on page 265.
- Configure PPPoE interfaces. See “[Example: Configuring PPPoE Interfaces](#)” on page 393.
- Configure PPPoE encapsulation on an Ethernet interface. See “[Example: Configuring PPPoE Encapsulation on an Ethernet Interface](#)” on page 400.

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

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

Confirm that the configuration is working properly.

- [Verifying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface on page 403](#)

[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.


Related Documentation

- [Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 401](#)

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 25 on page 401](#) shows a typical PPPoE over ATM session between a device and an access concentrator on an ADSL loop.

Figure 26: 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.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 389](#)
- [Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 402](#)

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

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 on page 405](#)
- [Overview on page 405](#)
- [Configuration on page 405](#)
- [Verification on page 406](#)

Requirements

Before you begin:

- Configure network interfaces. See [“Example: Creating an Ethernet Interface” on page 265](#).
- Configure PPPoE interfaces. See [“Example: Configuring PPPoE Interfaces” on page 393](#).
- Configure PPPoE encapsulation on an Ethernet interface. See [“Example: Configuring PPPoE Encapsulation on an Ethernet Interface” on page 400](#).

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

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

Confirm that the configuration is working properly.

- [Verifying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface on page 406](#)

[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.

Related Documentation • [Understanding PPPoE ATM-over-ADSL and ATM-over-SHDSL Interfaces on page 401](#)

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.

Related Documentation

- [Understanding Point-to-Point Protocol over Ethernet on page 389](#)
- [Example: Configuring CHAP Authentication on a PPPoE Interface on page 407](#)

Example: Configuring CHAP Authentication on a PPPoE Interface

This example shows how to configure CHAP authentication on a PPPoE interface.

- [Requirements on page 407](#)
- [Overview on page 407](#)
- [Configuration on page 408](#)
- [Verification on page 409](#)

Requirements

Before you begin:

- Configure an Ethernet interface. See [“Example: Creating an Ethernet Interface” on page 265](#).
- Configure a PPPoE interface. See [“Example: Configuring PPPoE Interfaces” on page 393](#).
- Configure PPPoE encapsulation on an ATM-over-ADSL interface. See [“Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface” on page 402](#).

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

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

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.

Related Documentation

- [Understanding CHAP Authentication on a PPPoE Interface on page 407](#)

Verifying Credit-Flow Control

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

```

- Related Documentation**
- [Understanding CHAP Authentication on a PPPoE Interface on page 407](#)
 - [Verifying Credit-Flow Control on page 409](#)

Verifying PPPoE Interfaces

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
```

- Related Documentation**
- [Understanding PPPoE Interfaces on page 392](#)
 - [Example: Configuring PPPoE Interfaces on page 393](#)

Verifying R2CP Interfaces

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		

- Related Documentation**
- [Understanding PPPoE Interfaces on page 392](#)
 - [Example: Configuring PPPoE Interfaces on page 393](#)

Displaying Statistics for PPPoE

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, mp1s:
  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::21f:12ff:fed2:2918
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: feee::5:1:1:0/126, Local: feee::5:1:1:1
```

- Related Documentation**
- [Understanding CHAP Authentication on a PPPoE Interface on page 407](#)
 - [Verifying Credit-Flow Control on page 409](#)

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

**Related
Documentation**

- [Understanding PPPoE Interfaces on page 392](#)
- [Example: Configuring PPPoE Interfaces on page 393](#)

CHAPTER 22

Configuring PPPoE-Based Radio-to-Router Protocol

- [PPPoE-Based Radio-to-Router Protocols Overview on page 415](#)
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 416](#)
- [Configuring PPPoE-Based Radio-to-Router Protocols on page 418](#)
- [Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 418](#)
- [Credit Flow Control for PPPoE on page 421](#)
- [PPPoE Credit-Based Flow Control Configuration on page 421](#)

PPPoE-Based Radio-to-Router Protocols Overview

Support for PPPoE-based radio-to-router protocols includes the following extensions to the PPPoE protocol:

- Messages that define how an external device provides the router with timely information about the quality of a link connection
- A flow control mechanism that indicates how much data the router can forward

The router uses the information provided in these PPPoE messages to dynamically adjust the interface speed. When OSPF is notified of this change, it adjusts the cost of the link and updates the routing tables accordingly.

The radio provides ground-to-ground or ground-to-air communications with like devices. When the radio picks up a signal from another device, it initiates a PPPoE session with a directly connected router. The PPPoE session encapsulates the packets that are relayed over a PPP link between the local and remote routers. The remote radio then forwards traffic over an independent PPPoE session between the remote radio and the router to which it is connected. The two routers exchange LCP and IPCP messages to configure the link and exchange OSPF messages to establish the network topology.

The router and radio are deployed in highly dynamic environments, such as moving vehicles. The quality of the radio link between the routers can vary significantly as a vehicle moves behind an obstruction. Each radio monitors the link every 50 milliseconds for changes in the link bandwidth, quality, and utilization. If any changes are detected, the radios announce the new set of metrics to the respective routers through a PPPoE Active Discovery Quality (PADQ) message, which is a nonstandard extension to the

PPPoE Discovery Protocol [RFC2516]. The router transforms these metrics into a bandwidth value for the PPP link and compares it to the value currently in use. When the router detects that the difference exceeds a user-specified threshold, it adjusts the speed of the PPP link. An event message notifies OSPF of the change, which then triggers OSPF to announce any resulting routing topology changes to its neighbors.

The PPPoE-based radio-to-router protocol notifies the router about neighbors joining or leaving the network and to create and maintain OSPF adjacencies over the dynamic links established between them. The costs assigned to these links are based on network conditions and flow control information sent by the radios. The calculations and requests to update interface speeds are performed by routines in a common library.

When PPPoE is used for applications, such as mobile radio, the radio links have variable bandwidth. So a mobile radio can function in a PPPoE environment, PPPoE messaging includes PADQ messages, which enable a link cost to be propagated to OSPF through the evaluation of various link quality metrics. The router uses information from these notifications along with user-configured parameters to calculate interface link costs that are used by the routing protocols.

A radio can send an optional PADQ at any time to query or report link quality metrics. When transmitting PPP streams over radio links, the quality of the link directly affects the throughput. The PADQ packet is used by the radio modem to report link metrics.

To support the credit-based flow control extensions described in RFC4938, PPPoE peers can also grant each other forwarding credits. The grantee can forward traffic to the peer only when it has a sufficient number of credits to do so. Credit-based forwarding allows both sides of the session to agree to use a non-default credit scaling factor during the PADR and PADS message exchange. Although this is used on both sides of the session, this feature provides the radio client with a flow control mechanism that throttles traffic by limiting the number of credits it grants to the router.

**Related
Documentation**

- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 416](#)

Understanding the PPPoE-Based Radio-to-Router Protocol

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.

For example, a high-band networking waveform (HNW) radio provides ground-to-ground or ground-to-air communications with like devices. When the HNW picks up a signal from another device, it initiates a PPPoE session with a directly connected device (router). The PPPoE session encapsulates the packets that are relayed over a PPP link between the local and remote devices. The remote radio then forwards traffic to a remote device using an independent PPPoE session. The two devices exchange Link Control Protocol

(LCP) and Internet Protocol Control Protocol (IPCP) messages to configure the link and exchange OSPF messages to establish the network topology.

Each HNW radio monitors the link every 50 milliseconds for changes in the link bandwidth, quality, and utilization. If any changes are detected, the radios announce the new set of metrics to the respective devices through a PPPoE Active Discovery Quality (PADQ) message, which is a nonstandard extension to the PPPoE Discovery Protocol (RFC 2516). The device transforms these metrics into a bandwidth value for the PPP link and compares it to the value currently in use. When the device detects that the difference exceeds a user-specified threshold, it adjusts the speed of the PPP link. OSPF is notified of the change and announces any resulting routing topology changes to its neighbors.

The CLI statement, **radio-router**, indicates that metrics announcements received on the interface will be processed by the device. When a PPPoE logical interface refers to this as an underlying interface, the device then processes incoming PADQ messages and uses information from the host's messages to control the flow of traffic and manage the speed of the link, resulting in a corresponding adjustment of the OSPF cost. If this option is not specified, then PADQ messages received over the underlying interface are ignored.

The following options are available within the **radio-router** configuration statement:

- **bandwidth, resource, latency, and quality**—These statements provide control over the weights used when transforming PADQ link metrics into an interface speed for the virtual link:
 - **bandwidth**—Weight of current (vs. maximum) data rate
 - **resource**—Resource weight
 - **latency**—Latency weight
 - **quality**—Relative link quality weight

All four weights accept values from 0 through 100. The default value for all four weights is 100.

- **credit**—This statement supports the credit-based flow control extensions described in RFC 4938. The statement enables PPPoE peers to grant each other forwarding credits. The grantee is then allowed to forward traffic to the peer only when it has a sufficient number of credits to do so. The subsequent credit interval statement controls how frequently the device generates credit announcement messages. The **interval** sub-statement, which controls the grant rate interval, accepts values from 1 through 60 seconds.
- **threshold**—This statement specifies how much of a difference is required between the calculated and the current interface speeds. The **threshold** value, expressed as a percentage, defaults to 10.

The following hierarchy provides another view of the **radio-router** configuration statements.

```
interfaces{
  interface-name {
    radio-router {
      bandwidth;
```

```
        credit {  
            interval;  
        }  
        latency;  
        quality;  
        resource;  
        threshold;  
    }  
}
```

- Related Documentation**
- [Understanding Point-to-Point Protocol over Ethernet on page 389](#)
 - [Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 418](#)

Configuring PPPoE-Based Radio-to-Router Protocols

To configure the PPPoE-based radio-to-router protocol:

1. Configure PPPoE encapsulation for an Ethernet interface.
2. Configure radio-router on the logical Ethernet interface.
3. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
4. Configure the operational mode as server.
5. (Optional) Identify the access concentrator by a unique name.
6. Specify how many seconds to wait before attempting to reconnect.
7. Provide a name for the type of service provided by the access concentrator.
8. Configure the maximum transmission unit (MTU) of the interface.
9. Configure the MTU size for the protocol family.
10. Disable the sending of keepalive messages on the logical interface.

- Related Documentation**
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 416](#)
 - [Example: Configuring the PPPoE-Based Radio-to-Router Protocol on page 418](#)

Example: Configuring the PPPoE-Based Radio-to-Router Protocol

This example shows how to configure the PPPoE-based radio-to-router protocol.

- [Requirements on page 419](#)
- [Overview on page 419](#)
- [Configuration on page 419](#)
- [Verification on page 420](#)

Requirements

Before you begin:

1. Configure network interfaces. See [“Example: Creating an Ethernet Interface” on page 265](#).
2. Configure PPPoE interfaces. See [“Example: Configuring PPPoE Interfaces” on page 393](#).
3. Configure PPPoE encapsulation on an Ethernet interface. See [“Example: Configuring PPPoE Encapsulation on an Ethernet Interface” on page 400](#).
4. Configure PPPoE encapsulation on an ATM-over-ADSL interface. See [“Example: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface” on page 402](#).
5. Configure CHAP authentication on a PPPoE interface. See [“Example: Configuring CHAP Authentication on a PPPoE Interface” on page 407](#).

Overview

In this example, you configure the ge-3/0/3 interface and set the bandwidth, resource, latency, and quality to **100**. You also set the threshold value to **10**, and then configure options on the logical interface.

Configuration

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.

```
[edit]
set interfaces ge-3/0/3 unit 1 radio-router bandwidth 100 resource 100 latency 100 quality
  100 threshold 10
set interfaces pp0 unit 1 pppoe-options underlying-interface ge-3/0/3 server
set interfaces pp0 unit 1 family inet unnumbered-address lo0.0 destination 192.168.1.2
set interfaces pp0 unit 1 family inet6 address lo0.0 destination fec0::1::2
```

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 the PPPoE-based radio-to-router protocol:

1. Enable the PPPoE-based radio-to-router protocol.

```
[edit]
user@host# edit interfaces ge-3/0/3 unit 1 radio-router
```

2. Set the interface speed for the virtual link.

```
[edit interfaces ge-3/0/3 unit 1 radio-router]
user@host# set bandwidth 100 resource 100 latency 100 quality 100
```

3. Set the calculated and current interface speeds, as a percentage.

```
[edit interfaces ge-3/0/3 unit 1 radio-router]
user@host# set threshold 10
```

4. Configure options on the logical interface.

```
[edit interfaces pp0 unit 1]
user@host# set pppoe-options underlying-interface ge-3/0/3
user@host# set pppoe-options server
user@host# set family inet unnumbered-address lo0.0 destination 192.168.1.2
user@host# set family inet6 address lo0.0 destination fec0:1:1::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.

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]
user@host# show interfaces ge-3/0/3 {
  unit 1
  radio-router {
    bandwidth 100;
    resource 100;
    latency 100;
    quality 100;
    threshold 10;
  }
}
...
pp0 {
  unit 1 {
    pppoe-options {
      underlying-interface ge-3/0/3;
      server;
    }
  }
  family inet {
    unnumbered-address lo0.0 destination 192.168.1.2;
  }
  family inet6;
}
}
```

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

Verification

Confirm that the configuration is working properly.

Verifying the PPPoE-based Radio-to-Router Protocol

Purpose Verify the PPPoE-Based radio-to-router protocol.

Action From operational mode, enter the **show interfaces** command.

Related Documentation

- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 416](#)

Credit Flow Control for PPPoE

To support the credit-based flow control extensions described in RFC4938, PPPoE peers can grant each other forwarding credits. The grantee is allowed to forward traffic to the peer only when it has a sufficient number of credits to do so. When credit-based forwarding is used on both sides of the session, the radio client can throttle traffic by limiting the number of credits it grants to the router.

The **interfaces** statement includes the **radio-router** attribute, which contains the parameters used for rate-based scheduling and OSPF link cost calculations. It also includes the **credit** attribute to indicate that credit-based packet scheduling is supported on the PPPoE interfaces that reference this underlying interface. Interfaces that set the **encapsulation** attribute support the PPPoE Active Discovery Grant (PADG) and PPPoE Active Discovery Credit (PADC) messages in the same way that the **radio-router** attribute provides active support for the PPPoE Active Discovery Quality (PADQ) message.

The **credit interval** parameter controls how frequently the router generates credit announcement messages. For PPPoE this corresponds to the interval between PADG credit announcements for each session.

Related Documentation

- [PPPoE-Based Radio-to-Router Protocols Overview on page 415](#)
- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 416](#)
- [Configuring PPPoE-Based Radio-to-Router Protocols on page 418](#)

PPPoE Credit-Based Flow Control Configuration

This example shows a PPPoE credit-based flow control configuration.

```
[edit interfaces ge-0/0/1]
unit 0 {
  encapsulation ppp-over-ether;
  radio-router {
    credit {
      interval 10;
    }
    bandwidth 80;
    threshold 5;
  }
}
```

}

**Related
Documentation**

- [Understanding the PPPoE-Based Radio-to-Router Protocol on page 416](#)
- [Configuring PPPoE-Based Radio-to-Router Protocols on page 418](#)

CHAPTER 23

Configuring R2CP Radio-to-Router Protocol

- [R2CP Radio-to-Router Protocol Overview on page 423](#)
- [Configuring the R2CP Radio-to-Router Protocol on page 424](#)

R2CP Radio-to-Router Protocol Overview

The Network Centric Waveform (NCW) radio-specific radio-to-router control protocol (R2CP) is similar to the PPPoE radio-to-router protocol. Both of these protocols exchange dynamic metric changes in the network that the routers use to update the OSPF topologies.

In radio-router topologies, the router connects to the radio over a Gigabit Ethernet link and the radio transmits packets over the radio frequency (RF) link. The radio periodically sends metrics to the router, which uses RF link characteristics and other data to inform the router on the shaping and OSPF link capacity. The router uses this information to shape the data traffic and provide the OSPF link cost for its SPF calculations. The radio functions like a Layer 2 switch and can only identify remote radio-router pairs using the Layer 2 MAC addresses. With R2CP the router receives metrics for each neighboring router, identified by the MAC address of the remote router. The R2CP daemon translates the MAC addresses to link the local IPv6 address and sends the metrics for each neighbor to OSPF. Processing these metrics is similar to the handling of PPPoE PADQ metrics. Unlike PPPoE, which is a point-to-point link, these R2CP neighbors are treated as nodes in a broadcast LAN.

You must configure each neighbor node with a per unit scheduler for CoS. The scheduler context defines the attributes of Junos class-of-service. To define CoS for each radio, you can configure virtual channels to limit traffic. You need to configure virtual channels for as many remote radio-router pairs as there are in the network. You configure virtual channels on a logical interface. Each virtual channel can be configured to have a set of eight queues with a scheduler and an optional shaper. When the radio initiates the session with a peer radio-router pair, a new session is created with the remote MAC address of the router and the VLAN over which the traffic flows. Junos OS chooses from the list of free virtual channels and assigns the remote MAC and the eight CoS queues and the scheduler to this remote MAC address. All traffic destined to this remote MAC address is subjected to the CoS that is defined in the virtual channel.

A virtual channel group is a collection of virtual channels. Each radio can have only one virtual channel group assigned uniquely. If you have more than one radio connected to the router, you must have one virtual channel group for each local radio-to-router pair. Although a virtual channel group is assigned to a logical interface, a virtual channel is not the same as a logical interface. The only features supported on a virtual channel are queuing, packet scheduling, and accounting. Rewrite rules and routing protocols apply to the entire logical interface.

All nodes in the R2CP network are in a broadcast LAN. The point-to-multipoint over LAN protocol supports advertising different bandwidth information for neighbors on a broadcast link. The network link is a point-to-multipoint link in the OSPFv3 link state database, which uses existing OSPF neighbor discovery to provide automatic discovery without configuration. It enables each node to advertise a different metric to every other node in the network to accurately represent the cost of communication. The **p2mp-over-lan** interface type under the OSPFv3 interface configuration enables you to configure the interface. OSPFv3 then uses LAN procedures for neighbor discovery and flooding, but represents the interface as point-to-multipoint in the link state database.

The interface type and router LSA are available under the following hierarchies:

[protocols ospf3 area *area-id* interface *interface-name*]

[routing-instances *routing-instances-name* protocols ospf3 area *area-id* interface *interface-name*]

For example:

```
protocols {
  ospf3 {
    area 0.0.0.0 {
      interface ge-0/0/2.0 {
        interface-type p2mp-over-lan;
      }
    }
  }
}
```

Related Documentation • [Configuring the R2CP Radio-to-Router Protocol on page 424](#)

Configuring the R2CP Radio-to-Router Protocol

To configure the R2CP protocol:

1. Configure the interfaces.

The following example creates four logical interfaces on ge-0/0/2, using unit 52 for R2CP control messages and units 101-193 for data traffic. The **per-unit-scheduler** statement is required for R2CP.

```
interfaces {
  ge-0/0/2 {
    per-unit-scheduler;
```



```

vlan-tagging;
unit 52 {
  vlan-id 52;
  family inet {
    address 52.1.1.1/24;
  }
}
unit 101 {
  vlan-id 101;
  family inet {
    address 101.1.1.1/24;
  }
}
unit 102 {
  vlan-id 102;
  family inet {
    address 102.1.1.1/24;
  }
}
unit 103 {
  vlan-id 103;
  family inet {
    address 103.1.1.1/24;
  }
}
}

```

2. Configure the R2CP protocol.

The following example configures `ge-0/0/2.52` as the interface for R2CP control messages, `vg1` as the virtual-channel group, and `ge-0/0/2.101-103` as data interfaces using the radio-interface statement.

```

protocols {
  r2cp {
    radio myRadio {
      interface ge-0/0/2.52;
      virtual-channel-group vg1;
      radio-interface ge-0/0/2.101;
      radio-interface ge-0/0/2.102;
      radio-interface ge-0/0/2.103;
    }
  }
}

```

3. Configure class of service.

The following example defines virtual-channels, their initial shaping-rates, and the virtual-channel-group to which they belong. It also makes the association between radio-interface interfaces and virtual-channel-group. In the class of service configuration, the **`vc-shared-scheduler`** configuration statement is required for each interface configured as a radio interface in the R2CP protocol configuration.

```

class-of-service {
  virtual-channels {
    vc1;
  }
}

```

```
    vc2;
    vc3;
    vc4;
}
virtual-channel-groups {
  vg1 {
    vc1 {
      scheduler-map sm;
      shaping-rate 15m;
      default;
    }
    vc2 {
      scheduler-map sm;
      shaping-rate 20m;
    }
    vc3 {
      scheduler-map sm;
      shaping-rate 20m;
    }
    vc4 {
      scheduler-map sm;
      shaping-rate 20m;
    }
  }
}
forwarding-classes {
  queue 0 DATA-queue;
}
interfaces {
  ge-0/0/2 {
    unit 101 {
      virtual-channel-group vg1;
      vc-shared-scheduler;
    }
    unit 102 {
      virtual-channel-group vg1;
      vc-shared-scheduler;
    }
    unit 103 {
      virtual-channel-group vg1;
      vc-shared-scheduler;
    }
  }
}
scheduler-maps {
  sm {
    forwarding-class DATA-queue scheduler sm-scheduler;
  }
}
schedulers {
  sm-scheduler {
    transmit-rate percent 20;
    buffer-size percent 20;
    priority low;
  }
}
```

}

Related Documentation

- [R2CP Radio-to-Router Protocol Overview on page 423](#)

PART 6

Configuring Link Services and Special Interfaces

- [Configuring Link Services Interfaces on page 431](#)
- [Configuring Link Fragmentation and Interleaving on page 455](#)
- [Configuring Class-of-Service on Link Services Interfaces on page 459](#)
- [Achieving Greater Bandwidth, Load Balancing, and Redundancy with Multilink Bundles on page 471](#)
- [Configuring Multilink Frame Relay on page 477](#)
- [Configuring Compressed Real-Time Transport Protocol on page 485](#)
- [Configuring Link Services Queuing Interface on page 489](#)
- [Understanding Special Interfaces on page 493](#)

Configuring Link Services Interfaces

- [Link Services Interfaces Overview on page 431](#)
- [Link Services Configuration Overview on page 437](#)
- [Verifying the Link Services Interface on page 438](#)
- [Troubleshooting the Link Services Interface on page 443](#)

Link Services Interfaces Overview

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 on page 432](#)
- [Link Services Exceptions on page 432](#)
- [Configuring Multiclass MLPPP on page 433](#)
- [Queuing with LFI on page 434](#)
- [Compressed Real-Time Transport Protocol Overview on page 435](#)

- [Configuring Fragmentation by Forwarding Class on page 435](#)
- [Configuring Link-Layer Overhead on page 437](#)

Services Available on a Link Services Interface

The link services interface is a logical interface available by default. [Table 38 on page 432](#) summarizes the services available on the interface.

Table 38: 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.	<ul style="list-style-type: none"> • Example: Configuring an MLPPP Bundle on page 472 • Example: Configuring Multilink Frame Relay FRF.15 on page 477 • Example: Configuring Multilink Frame Relay FRF.16 on page 481
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 455
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 435
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: <ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • Example: Configuring Interface Shaping Rates on page 468 • Configuring Fragmentation by Forwarding Class on page 435

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 434](#).
- 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-name multilink-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,

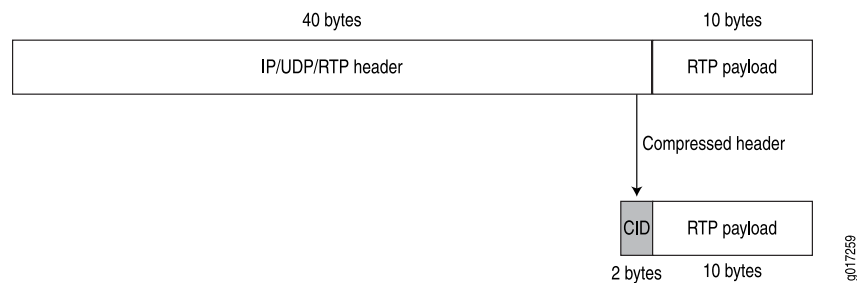
destination address, and the IP protocol of the packet. If the IP payload is TCP or UDP traffic, the hash also includes the source port and destination port. As a result of this hash algorithm, all traffic belonging to one traffic flow is queued to Q2 of one constituent link. This method of traffic delivery 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 27 on page 435 shows how CRTP compresses the RTP header in a voice packet by reducing a 40-byte header to a 2-byte header.

Figure 27: CRTP



You can configure CRTP with MLPPP or PPP logical interface encapsulation on link services interfaces. See [“Example: Configuring an MLPPP Bundle” on page 472](#).

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 455](#).

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.

Related Documentation

- [Link Services Configuration Overview on page 437](#)
- [Understanding the Internal Interface LSQ-0/0/0 Configuration on page 489](#)
- [Verifying the Link Services Interface on page 438](#)

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 3](#)

Plan how you are going to use the link services interface on your network. See [“Link Services Interfaces Overview” on page 431](#).

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 456](#).
2. Configure classifiers and forwarding classes. See [“Example: Defining Classifiers and Forwarding Classes” on page 460](#).
3. Configure scheduler maps. See [“Understanding How to Define and Apply Scheduler Maps” on page 463](#).
4. Configure interface shaping rates. See [“Example: Configuring Interface Shaping Rates” on page 468](#).
5. Configure an MLPPP bundle. See [“Example: Configuring an MLPPP Bundle” on page 472](#).
6. To configure MLFR, see [“Example: Configuring Multilink Frame Relay FRF.15” on page 477](#) or [“Example: Configuring Multilink Frame Relay FRF.16” on page 481](#).
7. To configure CRTP, see [“Example: Configuring the Compressed Real-Time Transport Protocol” on page 485](#).

Related Documentation

- [Link Services Interfaces Overview on page 431](#)
- [Understanding Multilink Frame Relay FRF.15 on page 477](#)
- [Understanding Multilink Frame Relay FRF.16 on page 480](#)
- [Understanding Compressed Real-Time Transport Protocol on page 485](#)
- [Understanding the Internal Interface LSQ-0/0/0 Configuration on page 489](#)
- [Verifying the Link Services Interface on page 438](#)

Verifying the Link Services Interface

Confirm that the configuration is working properly.

- [Verifying Link Services Interface Statistics on page 438](#)
- [Verifying Link Services CoS Configuration on page 441](#)

Verifying Link Services Interface Statistics

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 472](#). 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 472](#).
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.

See Also • [Link Services Interfaces Overview on page 431](#)

Verifying Link Services CoS Configuration

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 472](#).

```
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: classify_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.

- See Also**
- [Troubleshooting the Link Services Interface on page 443](#)
 - [Link Services Interfaces Overview on page 431](#)

Troubleshooting the Link Services Interface

To solve configuration problems on a link services interface:

- [Determine Which CoS Components Are Applied to the Constituent Links on page 443](#)
- [Determine What Causes Jitter and Latency on the Multilink Bundle on page 445](#)
- [Determine If LFI and Load Balancing Are Working Correctly on page 445](#)
- [Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device on page 452](#)

Determine Which CoS Components Are Applied to the Constituent Links

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 39 on page 443](#) shows the CoS components to be applied on a multilink bundle and its constituent links.

Table 39: 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 39: 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	<p>Apply scheduler maps on the multilink bundle and the constituent link as follows:</p> <ul style="list-style-type: none"> • 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.
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**
- [Link Services Interfaces Overview on page 431](#)
 - [Troubleshooting the Link Services Interface on page 443](#)
 - See the *Junos OS Class of Service Configuration Guide for Security Devices*

Determine What Causes Jitter and Latency on the Multilink Bundle

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).

- See Also**
- *RPM Overview*

Determine If LFI and Load Balancing Are Working Correctly

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
  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 28 on page 448 shows the overhead added to PPP and MLPPP headers.

Figure 28: 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 40 on page 448 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 40: 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
Data fragment (non-LFI) with short sequence	MLPPP	70 bytes	4 + 2 + 1 + 4 + 2 = 13 bytes	83 bytes

Table 40: PPP and MLPPP Encapsulation Overhead (continued)

Packet Type	Encapsulation	Initial Packet Size	Encapsulation Overhead	Packet Size after Encapsulation
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 41 on page 451](#) 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 41: 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.

- See Also**
- [Link Services Interfaces Overview on page 431](#)
 - [Troubleshooting the Link Services Interface on page 443](#)

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

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.

- See Also**
- [Link Services Interfaces Overview on page 431](#)

- [Troubleshooting the Link Services Interface on page 443](#)

Configuring Link Fragmentation and Interleaving

- [Understanding Link Fragmentation and Interleaving Configuration on page 455](#)
- [Example: Configuring Link Fragmentation and Interleaving on page 456](#)

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 29 on page 456](#) 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 29: 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.

Related Documentation

- [Link Services Interfaces Overview on page 431](#)
- [Example: Configuring Link Fragmentation and Interleaving on page 456](#)

Example: Configuring Link Fragmentation and Interleaving

This example shows how to configure LFI.

- [Requirements on page 456](#)
- [Overview on page 456](#)
- [Configuration on page 457](#)
- [Verification on page 457](#)

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

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

Verifying Link Fragmentation and Interleaving Configuration

Purpose Verify the LFI configuration.

Action From operational mode, enter the **show interfaces lsq-0/0/0** command.

Related Documentation

- [Understanding Link Fragmentation and Interleaving Configuration on page 455](#)
- [Troubleshooting the Link Services Interface on page 443](#)
- [Verifying the Link Services Interface on page 438](#)

Configuring Class-of-Service on Link Services Interfaces

- [Understanding How to Define Classifiers and Forwarding Classes on page 459](#)
- [Example: Defining Classifiers and Forwarding Classes on page 460](#)
- [Understanding How to Define and Apply Scheduler Maps on page 463](#)
- [Example: Configuring Scheduler Maps on page 465](#)
- [Understanding Interface Shaping Rates on page 468](#)
- [Example: Configuring Interface Shaping Rates on page 468](#)

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.
-

- Related Documentation**
- [Link Services Interfaces Overview on page 431](#)
 - [Example: Defining Classifiers and Forwarding Classes on page 460](#)

Example: Defining Classifiers and Forwarding Classes

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 on page 460](#)
- [Overview on page 460](#)
- [Configuration on page 460](#)
- [Verification on page 462](#)

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

- 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 *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

To confirm that the configuration is working properly, perform this task:

- [Verifying Classifiers and Forwarding Classes on page 463](#)

Verifying Classifiers and Forwarding Classes

Purpose	Verify the classifiers and the forwarding classes.
Action	From operational mode, enter the show class-of-service command.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Feature Support Reference for SRX Series and J Series Devices</i> • Understanding How to Define Classifiers and Forwarding Classes on page 459 • Link Services Interfaces Overview on page 431 • Troubleshooting the Link Services Interface on page 443 • Verifying the Link Services Interface on page 438

Understanding How to Define and Apply Scheduler Maps

Juniper Networks devices support per-unit scheduling **set class-of-service schedulers SO 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 42 on page 463](#) 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 42: 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.

**Related
Documentation**

- [Link Services Interfaces Overview on page 431](#)
- [Example: Configuring Scheduler Maps on page 465](#)
- [Example: Configuring an MLPPP Bundle on page 472](#)
- [Understanding Interface Shaping Rates on page 468](#)

Example: Configuring Scheduler Maps

This example shows how to configure scheduler maps to determine the transmission service level for each output queue.

- [Requirements on page 465](#)
- [Overview on page 465](#)
- [Configuration on page 465](#)
- [Verification on page 467](#)

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 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.

Configuration

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

```
set 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 *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

To confirm that the configuration is working properly, perform this task:

- [Verifying the Configuration of scheduler maps. on page 467](#)

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.

Related Documentation

- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)
- [Understanding How to Define and Apply Scheduler Maps on page 463](#)
- [Troubleshooting the Link Services Interface on page 443](#)
- [Verifying the Link Services Interface on page 438](#)

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.

Related Documentation

- [Link Services Interfaces Overview on page 431](#)
- [Example: Configuring Interface Shaping Rates on page 468](#)
- [Understanding How to Define and Apply Scheduler Maps on page 463](#)

Example: Configuring Interface Shaping Rates

This example shows how to configure interface shaping rates to control the maximum rate of traffic transmitted on an interface.

- [Requirements on page 468](#)
- [Overview on page 469](#)
- [Configuration on page 469](#)
- [Verification on page 469](#)

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” on page 565](#).

- 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 465](#).

Overview

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.

Configuration

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.

Related Documentation

- *Junos OS Feature Support Reference for SRX Series and J Series Devices*
- [Link Services Interfaces Overview on page 431](#)
- [Understanding Interface Shaping Rates on page 468](#)
- [Troubleshooting the Link Services Interface on page 443](#)
- [Verifying the Link Services Interface on page 438](#)

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

- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 471](#)
- [Example: Configuring an MLPPP Bundle on page 472](#)

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.

Related Documentation

- [Link Services Interfaces Overview on page 431](#)
- [Example: Configuring an MLPPP Bundle on page 472](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 477](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 481](#)

Example: Configuring an MLPPP Bundle

This example shows how to configure an MLPPP bundle to increase traffic bandwidth.

- [Requirements on page 472](#)
- [Overview on page 472](#)
- [Configuration on page 473](#)
- [Verification on page 475](#)

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 30 on page 473](#), 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 30: Configuring MLPPP and LFI on Serial Links



Configuration

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
```

- 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

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.

Related Documentation

- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 471](#)
- [Troubleshooting the Link Services Interface on page 443](#)
- [Verifying the Link Services Interface on page 438](#)

CHAPTER 28

Configuring Multilink Frame Relay

- [Understanding Multilink Frame Relay FRF.15 on page 477](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 477](#)
- [Understanding Multilink Frame Relay FRF.16 on page 480](#)
- [Example: Configuring Multilink Frame Relay FRF.16 on page 481](#)

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.

Related Documentation

- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 471](#)
- [Example: Configuring an MLPPP Bundle on page 472](#)
- [Link Services Interfaces Overview on page 431](#)
- [Example: Configuring Multilink Frame Relay FRF.15 on page 477](#)

Example: Configuring Multilink Frame Relay FRF.15

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 on page 478](#)
- [Overview on page 478](#)
- [Configuration on page 478](#)
- [Verification on page 480](#)

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

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 dlc1 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;
  dlc1 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

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.
Related Documentation	<ul style="list-style-type: none">• Understanding Multilink Frame Relay FRF.15 on page 477• Link Services Configuration Overview on page 437

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.

- Related Documentation**
- [Understanding MLPPP Bundles and Link Fragmentation and Interleaving \(LFI\) on Serial Links on page 471](#)
 - [Example: Configuring an MLPPP Bundle on page 472](#)
 - [Link Services Interfaces Overview on page 431](#)
 - [Example: Configuring Multilink Frame Relay FRF.16 on page 481](#)

Example: Configuring Multilink Frame Relay FRF.16

This example shows how to configure MLFR FRF.16 for additional bandwidth, load balancing, and redundancy.

- [Requirements on page 481](#)
- [Overview on page 481](#)
- [Configuration on page 481](#)
- [Verification on page 484](#)

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 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.

Configuration

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

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.

Related Documentation

- [Link Services Interfaces Overview on page 431](#)
- [Understanding Multilink Frame Relay FRF.16 on page 480](#)
- [Link Services Configuration Overview on page 437](#)

Configuring Compressed Real-Time Transport Protocol

- [Understanding Compressed Real-Time Transport Protocol on page 485](#)
- [Example: Configuring the Compressed Real-Time Transport Protocol on page 485](#)

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.

**Related
Documentation**

- [Link Services Interfaces Overview on page 431](#)
- [Example: Configuring the Compressed Real-Time Transport Protocol on page 485](#)

Example: Configuring the Compressed Real-Time Transport Protocol

This example shows how to configure CRTP to improve packet transmission, especially for time-sensitive voice packets.

- [Requirements on page 486](#)
- [Overview on page 486](#)
- [Configuration on page 486](#)
- [Verification on page 487](#)

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

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

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 on page 431](#)
- [Understanding Compressed Real-Time Transport Protocol on page 485](#)

CHAPTER 30

Configuring Link Services Queuing Interface

- [Understanding the Internal Interface LSQ-0/0/0 Configuration on page 489](#)
- [Example: Upgrading from ls-0/0/0 to lsq-0/0/0 for Multilink Services on page 489](#)

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.

Related Documentation

- [Link Services Interfaces Overview on page 431](#)
- [Example: Upgrading from ls-0/0/0 to lsq-0/0/0 for Multilink Services on page 489](#)

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

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 on page 489](#)
- [Overview on page 490](#)
- [Configuration on page 490](#)
- [Verification on page 492](#)

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

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

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. |
| Related Documentation | <ul style="list-style-type: none">• Link Services Interfaces Overview on page 431• Understanding the Internal Interface LSQ-0/0/0 Configuration on page 489 |

Understanding Special Interfaces

- [Understanding Management Interfaces on page 493](#)
- [Understanding the Discard Interface on page 494](#)
- [Understanding the Loopback Interface on page 494](#)
- [Configuring a Loopback Interface on page 495](#)

Understanding 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.

- Related Documentation**
- [Understanding Interfaces on page 3](#)
 - [Understanding the Discard Interface on page 494](#)
 - [Understanding the Loopback Interface on page 494](#)

Understanding the 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 on page 3](#)
 - [Understanding Management Interfaces on page 493](#)
 - [Understanding the Loopback Interface on page 494](#)

Understanding the Loopback Interface

The loopback address (**lo0**) has several uses, depending on the particular Junos feature being configured. It can perform the following functions:

- **Device identification**—The loopback interface is used to identify the device. While any interface address can be used 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's configuration or operation.

- **Routing information**—The loopback address is used by protocols such as OSPF to determine protocol-specific properties for the device or network. Further, some commands such as **ping mpls** require a loopback address to function correctly.
- **Packet filtering**—Stateless firewall filters can be applied to the loopback address to filter packets originating from, or destined for, the Routing Engine.

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**.

The device also includes an internal loopback address (**lo0.16384**). The internal loopback address is a particular instance of the loopback address with the logical unit number 16384. Junos OS creates the loopback interface for the internal routing instance. This interface prevents any filter on **lo0.0** from disrupting internal traffic.

- Related Documentation**
- [Configuring a Loopback Interface on page 495](#)
 - [Understanding Interfaces on page 3](#)
 - [Understanding Management Interfaces on page 493](#)
 - [Understanding the Discard Interface on page 494](#)

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 master 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
```

**Related
Documentation**

- [Understanding the Loopback Interface on page 494](#)

PART 7

Configuring LTE Interfaces

- [Configuring LTE Interfaces on page 499](#)

CHAPTER 32

Configuring LTE Interfaces

- [LTE Mini-PIM Overview on page 499](#)
- [LTE Mini-PIM Configuration Overview on page 502](#)
- [Configuring the LTE Mini-PIM as the Primary Interface on page 503](#)
- [Configuring the LTE Mini-PIM as a Backup Interface on page 505](#)
- [Configuring the LTE Interface as a Dial-on-Demand Interface on page 507](#)

LTE Mini-PIM Overview

The LTE Mini-Physical Interface Module (Mini-PIM) provides wireless WAN support on the SRX320, SRX340, SRX345, and SRX550M (High Memory) Services Gateways. The LTE Mini-PIM operates on both 3G and 4G networks. [Table 43 on page 499](#) provides a summary of the different models of the Mini-PIM.

Table 43: LTE Mini-PIM Models

Model	Mode	Operating Region	Frequency Band
SRX-MP-LTE-AE	<ul style="list-style-type: none">• LTE• HSPA+	<ul style="list-style-type: none">• North America• European Union	For LTE: <ul style="list-style-type: none">• Bands 1 through 5, 7, 8, 12, 13, 20, 25, 26, 29, 30, and 41 For 3G (HSPA+): <ul style="list-style-type: none">• Bands 1 through 5, and 8
SRX-MP-LTE-AA	<ul style="list-style-type: none">• LTE• HSPA+	<ul style="list-style-type: none">• Asia• Australia	For LTE: <ul style="list-style-type: none">• Bands 1,3, 5, 7, 8, 18, 19, 21, 28, 38, 39, 40, and 41 For 3G (HSPA+): <ul style="list-style-type: none">• Bands 1, 5, 6, 8, 9, and 19

- [Supported Features on page 500](#)
- [Understanding the LTE Physical Interface on page 501](#)
- [Understanding the LTE Logical Interface on page 501](#)

Supported Features

The LTE Mini-PIM supports the following features:

- Automatic switchover between service providers through dual SIMs—The Mini-PIM supports up to two Subscriber Identity Module (SIM) cards. Dual SIM cards allow connectivity to two different ISP networks and provide a failover mechanism when the current active network fails. Each SIM card is associated with a profile, which is used to connect to the network.
- Multiple service provider and access point name (APN) profiles—You can configure up to 16 profiles for each SIM, although only one profile can be active at a time. The LTE Mini-PIM supports two SIM cards and so you can configure a total of 32 profiles.
- LTE carrier aggregation—Carrier aggregation expands the LTE bandwidth by combining secondary bands, which results in increased capacity and network efficiency.
- SIM security functions—The Mini-PIM supports security functions such as SIM lock and unlock, and PIN change.
- Always-on, dial-on-demand, and backup modes—The Mini-PIM can be configured in three modes:
 - Always-on—The Mini-PIM connects to the 3G/4G network after booting. The connection is always maintained, as long as there are no network or connectivity problems.
 - Dial-on-demand—The Mini-PIM initiates a connection when it receives interesting traffic. You define interesting traffic using the dialer filter. To configure dial-on-demand using a dialer filter, you first configure the dialer filter and then apply the filter to the dialer interface.



NOTE: The dial-on-demand mode is supported only if the LTE mini-PIM is configured as a primary interface.

- Backup—The Mini-PIM connects to the 3G/4G network when the primary connection fails.
- Primary and backup interface—You can configure the LTE Mini-PIM either as a primary interface or as a backup interface.

When configured as the primary interface, the LTE Mini-PIM supports both the Always-on and Dial-on-demand modes.

When configured as the backup interface, the LTE Mini-PIM connects to the network only when the primary interface fails.

- Over-the-air upgrade for modem firmware—Over-the-air (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.



NOTE: OTA upgrade is disabled by default.

Understanding the LTE Physical Interface

The physical interface for the 4G LTE Mini-PIM uses the name *cl-slot number/0/0*, where *slot number* identifies the slot on the services gateway in which you insert the Mini-PIM. For example, *cl-1/0/0*. The Mini-PIM can be inserted in any of the Mini-PIM slots on the SRX320, SRX340, SRX345, and SRX550M Services Gateways. You configure the following properties on the physical interface:

- 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)

Understanding the LTE Logical Interface

The dialer interface, *dl0*, is a logical interface, which is used to trigger calls. When traffic is sent to the *dl0* interface, it enables the physical interface in the dialer pool and places calls through the physical interface. The dialer interface can perform backup and dialer filter functions. You can configure the dialer interface to operate in any one of the following ways:

- 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” on page 503](#).
- Backup interface for the primary WAN connection—The dialer interface is activated only when the primary connection fails. For more information, see [“Configuring the LTE Mini-PIM as a Backup Interface” on page 505](#).
- Dial-on-demand—The dialer interface activates the connection to the wireless network only when it receives interesting traffic. For more information, see [“Configuring the LTE Interface as a Dial-on-Demand Interface” on page 507](#).

The following rules apply when you configure 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.
- Dial string (destination number to be dialed).

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.
- 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.

Class of Service on the Dialer Interface

The dialer interface has limited bandwidth, which can lead to traffic congestion. Starting with Junos OS Release 15.1X49-D100, the dialer interface supports the configuration of 4G LTE dialer interface Class of Service (CoS) parameters on SRX320, SRX340, SRX345, and SRX550M devices. The dialer interface supports the following CoS parameters:

- Behavior aggregate and multifield classifiers
- Policers
- Shapers
- Schedulers
- Rewrite rules



NOTE: The dialer interface (dl0) supports scheduling only at the physical interface queue level. As this interface does not support shaping at the logical interface level, per-unit scheduling is not supported on the dialer interface.

See *Class of Service Feature Guide for Security Devices* for information on configuring these parameters.

Related Documentation

- [LTE Mini-PIM Configuration Overview on page 502](#)

LTE Mini-PIM Configuration Overview

The configuration process for the LTE Mini-PIM includes the following tasks:

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. Establish an account with a cellular network service provider. Contact your service provider for more information.
3. Gather 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))

4. Install the LTE Mini-PIM.
5. Configure the LTE Mini-PIM. See:
 - [Configuring the LTE Mini-PIM as the Primary Interface on page 503](#)
 - [Configuring the LTE Mini-PIM as a Backup Interface on page 505](#)
 - [Configuring the LTE Interface as a Dial-on-Demand Interface on page 507](#)

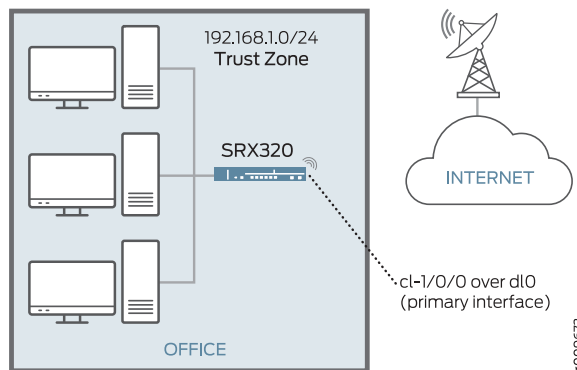
Configuring the LTE Mini-PIM as the Primary Interface

Figure 31 on page 503 illustrates a scenario where the LTE Mini-PIM is installed on a SRX320 Services Gateway and functions as the primary interface. This procedure assumes that the LTE Mini-PIM is installed in slot 1 on the SRX320 Services Gateway.



NOTE: The LTE Mini-PIM can be installed in any of the Mini-PIM slots on the SRX320, SRX340, SRX345, and SRX550M Services Gateways.

Figure 31: LTE Mini-PIM Used as a Primary Interface



Before you begin the procedure, ensure that dl0.0 is not configured as a backup. If dl0.0 is configured as a backup option 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

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. 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.

```
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
```



NOTE: *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
```



NOTE: 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 as well.

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
```




NOTE: If the LTE Mini-PIM gets an IP address with a mask of /32 from the service provider, the user has to 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.

Related Documentation

- [Configuring the LTE Interface as a Dial-on-Demand Interface on page 507](#)

Configuring the LTE Mini-PIM as a Backup Interface


You can 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.

[Figure 31 on page 503](#) illustrates a scenario where the LTE Mini-PIM is installed on a SRX320 Services Gateway and functions as a backup interface. The ge-0/0/1 port is connected to the Internet and functions as the primary interface.



NOTE: The LTE Mini-PIM can be installed in any of the Mini-PIM slots on the SRX320, SRX340, SRX345, and SRX550M Services Gateways. In this scenario, the Mini-PIM is installed on slot 1.

Figure 32: LTE Mini-PIM Used 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. 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.

```
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
```



NOTE: *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
```



NOTE: 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 as well.

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
```



NOTE: The `activation-delay` and `deactivation-delay` command-line options can be used to avoid 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 the activation of the dialer interface. Similarly, the deactivation delay controls the time between the recovery of the primary interface and the deactivation of the backup interface.

**Related
Documentation**

- [Configuring the LTE Mini-PIM as the Primary Interface on page 503](#)
- [Configuring the LTE Interface as a Dial-on-Demand Interface on page 507](#)

Configuring the LTE Interface as a Dial-on-Demand Interface

When the LTE interface is configured as a primary interface, it can function either in always-on mode 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, the dialer interface is enabled 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.



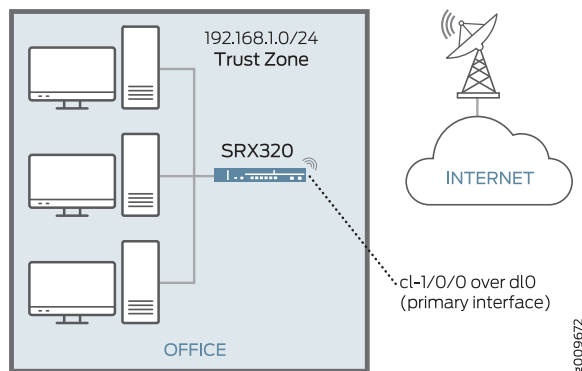
NOTE: The dial-on-demand mode is supported only if the LTE Mini-PIM is configured as a primary interface.

[Figure 31 on page 503](#) illustrates a scenario where the LTE Mini-PIM is installed on a SRX320 Services Gateway and functions as the primary interface. This procedure assumes that the LTE Mini-PIM is installed in slot 1 on the SRX320 Services Gateway.



NOTE: The LTE Mini-PIM can be installed in any of the Mini-PIM slots on the SRX320, SRX340, SRX345, and SRX550M Services Gateways.

Figure 33: LTE Mini-PIM Used 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
```



NOTE: Optionally, you can configure the `idle-timeout` value, which determines the duration for which the connection will remain enabled 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. 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.

```
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
```



NOTE: *sim-slot-number* is the slot on the Mini-PIM in which the SIM card is inserted.

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
```



NOTE: 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 as well.

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
```

Related Documentation

- [Configuring the LTE Mini-PIM as the Primary Interface on page 503](#)

PART 8

Configuring Modem Interfaces

- [Configuring 3G Wireless Modems for WAN Connections on page 513](#)
- [Configuring CDMA EV-DO Modem Cards on page 529](#)
- [Configuring USB Modems for Dial Backup on page 537](#)
- [Configuring DOCSIS Mini-PIM Interfaces on page 557](#)
- [Configuring Serial Interfaces on page 565](#)

CHAPTER 33

Configuring 3G Wireless Modems for WAN Connections

- [3G Wireless Modem Overview on page 513](#)
- [3G Wireless Modem Configuration Overview on page 514](#)
- [Understanding the Dialer Interface on page 516](#)
- [Example: Configuring the Dialer Interface on page 518](#)
- [Understanding the 3G Wireless Modem Physical Interface on page 523](#)
- [Example: Configuring the 3G Wireless Modem Interface on page 524](#)
- [Understanding the GSM Profile on page 525](#)
- [Example: Configuring the GSM Profile on page 526](#)


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 34 on page 514](#) 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 34: Wireless WAN Connections for Branch Offices



Related Documentation

- [3G Wireless Modem Configuration Overview on page 514](#)

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 518](#).
2. Configure the 3G wireless modem interface. See [“Example: Configuring the 3G Wireless Modem Interface” on page 524](#).
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.

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [Understanding the GSM Profile on page 525](#)
- [Unlocking the GSM 3G Wireless Modem on page 535](#)
- [Understanding Account Activation for CDMA EV-DO Modem Cards on page 529](#)

Understanding the Dialer Interface

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, and SRX550HM devices.

This topic contains the following sections:

- [Dialer Interface Configuration Rules on page 516](#)
- [Dialer Interface Authentication Support for GSM HSDPA 3G Wireless Modems on page 517](#)
- [Dialer Interface Functions on page 517](#)
- [Dialer Interface Operating Parameters on page 517](#)

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 **Isq-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.

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [3G Wireless Modem Configuration Overview on page 514](#)
- [Example: Configuring the Dialer Interface on page 518](#)

Example: Configuring the Dialer Interface

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, and SRX550HM devices.

- [Requirements on page 518](#)
- [Overview on page 518](#)
- [Configuration on page 518](#)
- [Verification on page 523](#)

Requirements

Before you begin, install your SRX Series device and establish basic connectivity for your device. See [“3G Wireless Modem Configuration Overview” on page 514](#).

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

- [Configuring a Dialer Interface on page 518](#)
- [Configuring PAP on the Dialer Interface on page 519](#)
- [Configuring CHAP on the Dialer Interface on page 520](#)
- [Configuring the Dialer Interface as a Backup WAN Connection on page 521](#)
- [Configuring Dialer Watch for the 3G Wireless Modem Interface on page 521](#)
- [Configuring a Dialer Filter for the 3G Wireless Modem Interface on page 522](#)

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

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.

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [3G Wireless Modem Configuration Overview on page 514](#)
- [Understanding the Dialer Interface on page 516](#)

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, 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.

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [3G Wireless Modem Configuration Overview on page 514](#)
- [Example: Configuring the 3G Wireless Modem Interface on page 524](#)

Example: Configuring the 3G Wireless Modem Interface

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, and SRX550HM devices.

- [Requirements on page 524](#)
- [Overview on page 524](#)
- [Configuration on page 524](#)
- [Verification on page 525](#)

Requirements

Before you begin, configure a dialer interface. See “[Example: Configuring the Dialer Interface](#)” on page 518.

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

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.

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [3G Wireless Modem Configuration Overview on page 514](#)
- [Understanding the 3G Wireless Modem Physical Interface on page 523](#)

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, and SRX550HM devices.

- Related Documentation**
- [3G Wireless Modem Overview on page 513](#)
 - [3G Wireless Modem Configuration Overview on page 514](#)
 - [Example: Configuring the GSM Profile on page 526](#)

Example: Configuring the GSM Profile

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, and SRX550HM devices.

-
- [Requirements on page 526](#)
 - [Overview on page 526](#)
 - [Configuration on page 526](#)
 - [Verification on page 527](#)

Requirements

Before you begin:

- Configure a dialer interface. See [“Example: Configuring the Dialer Interface” on page 518](#)
- Configure the 3G wireless modem interface. See [“Example: Configuring the 3G Wireless Modem Interface” on page 524](#).

Overview

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`.

Configuration

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.

- Related Documentation**
- [3G Wireless Modem Overview on page 513](#)
 - [3G Wireless Modem Configuration Overview on page 514](#)
 - [Understanding the GSM Profile on page 525](#)

Configuring CDMA EV-DO Modem Cards

- [Understanding Account Activation for CDMA EV-DO Modem Cards on page 529](#)
- [Activating the CDMA EV-DO Modem Card with IOTA Provisioning on page 531](#)
- [Activating the CDMA EV-DO Modem Card with OTASP Provisioning on page 532](#)
- [Activating the CDMA EV-DO Modem Card Manually on page 533](#)
- [Unlocking the GSM 3G Wireless Modem on page 535](#)

Understanding Account Activation for CDMA EV-DO Modem Cards

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, 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 \(ESN\) on page 529](#)
- [Account Activation Modes on page 530](#)

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.

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [3G Wireless Modem Configuration Overview on page 514](#)
- [Example: Configuring the GSM Profile on page 526](#)

Activating the CDMA EV-DO Modem Card with IOTA Provisioning

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, and SRX550HM devices.

Before you begin, activate the CDMA EV-DO 3G wireless modem card. See [“Understanding Account Activation for CDMA EV-DO Modem Cards” on page 529](#).

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
```

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [Activating the CDMA EV-DO Modem Card with OTASP Provisioning on page 532](#)
- [Activating the CDMA EV-DO Modem Card Manually on page 533](#)

Activating the CDMA EV-DO Modem Card with OTASP Provisioning

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, and SRX550HM devices.

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.

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 *2286*, 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
```

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [Understanding Account Activation for CDMA EV-DO Modem Cards on page 529](#)
- [Activating the CDMA EV-DO Modem Card Manually on page 533](#)
- [Activating the CDMA EV-DO Modem Card with IOTA Provisioning on page 531](#)

Activating the CDMA EV-DO Modem Card Manually

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, and SRX550HM devices.

Before you begin, the service provider must activate your account before you can activate the CDMA EV-DO 3G wireless modem card.

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
```

Related Documentation

- [3G Wireless Modem Overview on page 513](#)
- [Understanding Account Activation for CDMA EV-DO Modem Cards on page 529](#)
- [Activating the CDMA EV-DO Modem Card with OTASP Provisioning on page 532](#)
- [Activating the CDMA EV-DO Modem Card with IOTA Provisioning on page 531](#)

Unlocking the GSM 3G Wireless Modem

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 no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Before you begin, obtain the PIN from the service provider.

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 Unlocking 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.

**Related
Documentation**

- [3G Wireless Modem Overview on page 513](#)
- [3G Wireless Modem Configuration Overview on page 514](#)
- [Understanding the Dialer Interface on page 516](#)
- [Understanding the 3G Wireless Modem Physical Interface on page 523](#)
- [Understanding the GSM Profile on page 525](#)

Configuring USB Modems for Dial Backup

- [USB Modem Interface Overview on page 537](#)
- [USB Modem Configuration Overview on page 540](#)
- [Example: Configuring a USB Modem Interface on page 542](#)
- [Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup on page 545](#)
- [Example: Configuring a Dialer Interface for USB Modem Dial-In on page 551](#)
- [Example: Configuring PAP on Dialer Interfaces on page 553](#)
- [Example: Configuring CHAP on Dialer Interfaces on page 554](#)

USB Modem Interface Overview

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, 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 **umdn0**. 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, **dln**, 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 44 on page 539](#) describes the commands. For more information about these commands, see the documentation for your modem.

Table 44: 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.)

**Related
Documentation**

- [USB Modem Configuration Overview on page 540](#)
- [Example: Configuring a USB Modem Interface on page 542](#)
- [Example: Configuring a Dialer Interface for USB Modem Dial-In on page 551](#)

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.

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 45 on page 541](#) for a summarized description of the procedure.

Table 45: 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" on page 542 .
	Configure the dialer interface dl0 on the branch office router using one of the following backup methods: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 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 545. To configure a dialer filter on dl0, see "Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup" on page 545. To configure a dialer watch on dl0, see "Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup" on page 545.
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" on page 551 .

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 46 on page 542](#) for a list of available incoming map options.

Table 46: Incoming Map Options

Option	Description
accept-all	<p>Dialer interface accepts all incoming calls.</p> <p>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.</p>
caller	<p>Dialer interface accepts calls from a specific caller ID. You can configure a maximum of 15 caller IDs per dialer interface.</p> <p>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.</p>

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.

- Related Documentation**
- [USB Modem Interface Overview on page 537](#)
 - [Example: Configuring a USB Modem Interface on page 542](#)

Example: Configuring a USB Modem Interface

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 on page 543](#)
- [Overview on page 543](#)
- [Configuration on page 543](#)
- [Verification on page 544](#)

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

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 "ATSO=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

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 : ATSO=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

```

- Related Documentation**
- [USB Modem Configuration Overview on page 540](#)
 - [USB Modem Interface Overview on page 537](#)
 - [Example: Configuring a Dialer Interface for USB Modem Dial-In on page 551](#)

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

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, and SRX550HM devices.

- [Requirements on page 545](#)
- [Overview on page 545](#)
- [Configuration on page 545](#)
- [Verification on page 551](#)

Requirements

Before you begin, configure a USB modem for the device. See [“Example: Configuring a USB Modem Interface” on page 542](#).

Overview

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.

Configuration

- [Configuring a Dialer Interface for USB Modem Dial Backup on page 546](#)
- [Configuring a Dial Backup for a USB Modem Connection on page 547](#)

- [Configuring a Dialer Filter for USB Modem Dial Backup on page 548](#)
- [Configuring a Dialer Watch for USB Modem Dial Backup on page 550](#)

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

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.

Related Documentation

- [USB Modem Configuration Overview on page 540](#)
- [Example: Configuring a Dialer Interface for USB Modem Dial-In on page 551](#)
- [Example: Configuring PAP on Dialer Interfaces on page 553](#)
- [Example: Configuring CHAP on Dialer Interfaces on page 554](#)

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

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 on page 551](#)
- [Overview on page 552](#)
- [Configuration on page 552](#)
- [Verification on page 553](#)

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 d10.

Configuration

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 d10 unit 0 dialer-options incoming-map caller 4085550115
```

Step-by-Step Procedure

To configure a dialer interface for USB modem dial-in:

1. Select a dialer interface.

```
[edit]  
user@host# edit interfaces d10
```


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.

- Related Documentation**
- [USB Modem Configuration Overview on page 540](#)
 - [Example: Configuring a USB Modem Interface on page 542](#)

Example: Configuring PAP on Dialer Interfaces

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, and SRX550HM devices.

- [Requirements on page 553](#)
- [Overview on page 553](#)
- [Configuration on page 553](#)
- [Verification on page 554](#)

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

- 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.

Related Documentation

- [USB Modem Configuration Overview on page 540](#)
- [Example: Configuring a USB Modem Interface on page 542](#)
- [Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup on page 545](#)
- [Example: Configuring a Dialer Interface for USB Modem Dial-In on page 551](#)
- [Example: Configuring CHAP on Dialer Interfaces on page 554](#)

Example: Configuring CHAP on Dialer Interfaces

This example shows how to configure CHAP on dialer interfaces for authentication.



NOTE: Configuring CHAP on dialer interfaces for authentication is no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

- [Requirements on page 555](#)
- [Overview on page 555](#)
- [Configuration on page 555](#)
- [Verification on page 555](#)

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 d10. Finally, you enable CHAP on a dialer interface and specify a unique profile name containing a client list and access parameters.

Configuration

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 d10 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 d10** command.

Related Documentation

- [USB Modem Configuration Overview on page 540](#)
- [Example: Configuring a USB Modem Interface on page 542](#)
- [Example: Configuring Dialer Interfaces and Backup Methods for USB Modem Dial Backup on page 545](#)
- [Example: Configuring a Dialer Interface for USB Modem Dial-In on page 551](#)
- [Example: Configuring PAP on Dialer Interfaces on page 553](#)

Configuring DOCSIS Mini-PIM Interfaces

- [DOCSIS Mini-PIM Interface Overview on page 557](#)
- [Software Features Supported on DOCSIS Mini-PIMs on page 559](#)
- [Example: Configuring the DOCSIS Mini-PIM Interfaces on page 560](#)

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, 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.



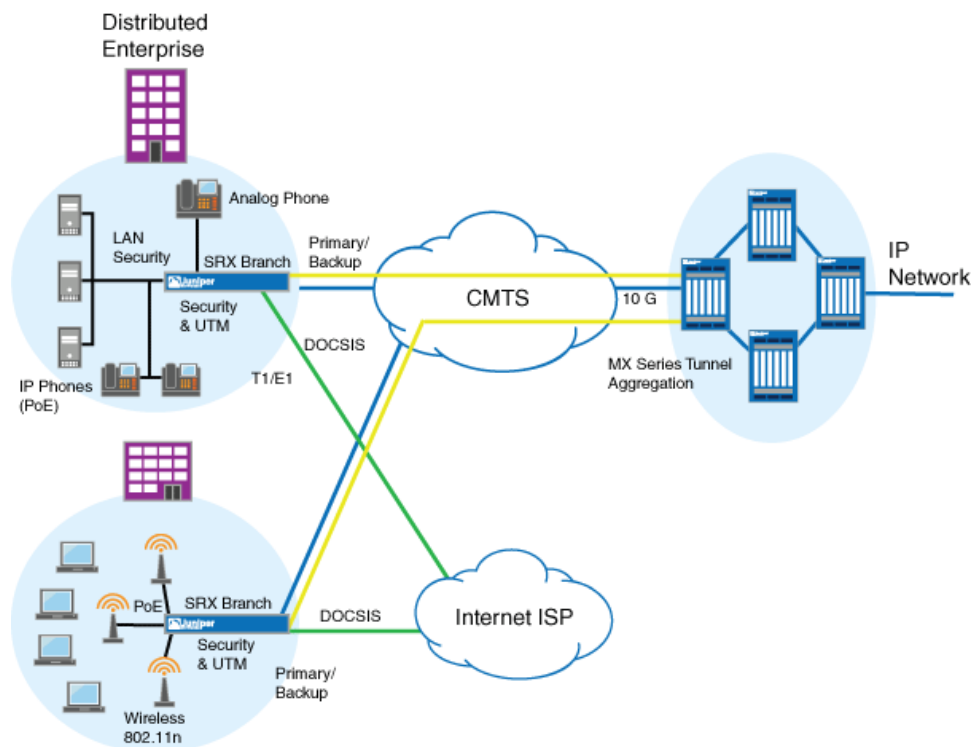
NOTE: The following interface trace options are supported:

- **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 35 on page 558](#) shows a typical use for this Mini-PIM in an MSO network.

Figure 35: Typical DOCSIS End-to-End Connectivity Diagram



Related Documentation

- [Software Features Supported on DOCSIS Mini-PIMs on page 559](#)
- [Example: Configuring the DOCSIS Mini-PIM Interfaces on page 560](#)

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 47 on page 559 lists the software features supported on DOCSIS Mini-PIMs.

Table 47: Software Features Supported on DOCSIS Mini-PIMs

Software Feature	Description
DHCP and DHCPv6 clients	<p>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:</p> <ul style="list-style-type: none"> 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	<p>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.</p> <p>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.</p>
SNMP support	<p>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:</p> <ul style="list-style-type: none"> NAT support Dying gasp support Back pressure information
MAC address	<p>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.</p>
Transparent bridging	<p>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.</p>

Release History Table

Release	Description
15.1X49-D10	DOCSIS Mini-PIM interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Related Documentation

- [DOCSIS Mini-PIM Interface Overview on page 557](#)
- [Example: Configuring the DOCSIS Mini-PIM Interfaces on page 560](#)

Example: Configuring the DOCSIS Mini-PIM Interfaces

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, and SRX550HM devices.

- [Requirements on page 560](#)
- [Overview on page 560](#)
- [Configuration on page 560](#)
- [Verification on page 562](#)

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 265](#).

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

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

Confirm that the configuration is working properly.

- [Verifying the DOCSIS Interface Properties on page 562](#)

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:
  Total octets                Receive      Transmit
  Total packets               1935          2036
  CRC/Align errors            8
  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
  Output bytes  : 866
  Input packets: 2
  Output packets: 4
  Packet Forwarding Engine configuration:
    Destination slot: 1
    Direction : Output
    CoS transmit queue
    Limit
    Bandwidth
    Buffer Priority
    %
    bps
    %
    usec
    low
  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.

**Related
Documentation**

- [DOCSIS Mini-PIM Interface Overview on page 557](#)
- [Software Features Supported on DOCSIS Mini-PIMs on page 559](#)

CHAPTER 37

Configuring Serial Interfaces

- [Serial Interfaces Overview on page 565](#)
- [Example: Configuring a Serial Interface on page 571](#)
- [Example: Deleting a Serial Interface on page 574](#)
- [Understanding the 8-Port Synchronous Serial GPIM on page 575](#)
- [Example: Configuring an 8-Port Synchronous Serial GPIM in Back-to-Back SRX650 Services Gateways on page 577](#)

Serial Interfaces Overview

Serial links are simple, bidirectional links that require very few control signals. In a basic serial setup, data communications equipment (DCE) installed in a user's premises is responsible for establishing, maintaining, and terminating a connection. A modem is a typical DCE device.

A serial cable connects the DCE to a telephony network where, ultimately, a link is established with data terminal equipment (DTE). DTE is typically where a serial link terminates.

The distinction between DCE and DTE is important because it affects the cable pinouts on a serial cable. A DCE cable uses a female 9-pin or 25-pin connector, and a DTE cable uses a male 9-pin or 25-pin connector, and .

To form a serial link, the cables are connected to each other. However, if the pins are identical, each side's transmit and receive lines are connected, which makes data transport impossible. To address this problem, each cable is connected to a null modem cable, which crosses the transmit and receive lines in the cable.

This section includes the following topics:

- [Serial Transmissions on page 566](#)
- [Signal Polarity on page 567](#)
- [Serial Clocking Modes on page 567](#)
- [Serial Line Protocols on page 568](#)

Serial Transmissions

In basic serial communications, nine signals are critical to the transmission. Each signal is associated with a pin in either the 9-pin or 25-pin connector. [Table 48 on page 566](#) lists and defines serial signals and their sources.

Table 48: Serial Transmission Signals

Signal Name	Definition	Signal Source
TD	Transmitted data	DTE
RD	Received data	DCE
RTS	Request to send	DTE
CTS	Clear to send	DCE
DSR	Data set ready	DCE
Signal Ground	Grounding signal	–
CD	Carrier detect	–
DTR	Data terminal ready	DTE
RI	Ring indicator	–

When a serial connection is made, a serial line protocol—such as EIA-530, X.21, RS-422/449, RS-232, or V.35—begins controlling the transmission of signals across the line as follows:

1. The DCE transmits a DSR signal to the DTE, which responds with a DTR signal. After this handshake, the link is established and traffic can pass.
2. When the DTE device is ready to receive data, it sets its RTS signal to a marked state (all 1s) to indicate to the DCE that it can transmit data. (If the DTE is not able to receive data—because of buffer conditions, for example—it sets the RTS signal to all 0s.)
3. When the DCE device is ready to receive data, it sets its CTS signal to a marked state to indicate to the DTE that it can transmit data. (If the DCE is not able to receive data, it sets the CTS signal to all 0s.)
4. When the negotiation to send information has taken place, data is transmitted across the transmitted data (TD) and received data (RD) lines:
 - TD line—Line through which data from a DTE device is transmitted to a DCE device
 - RD line—Line through which data from a DCE device is transmitted to a DTE device

The name of the wire does not indicate the direction of data flow.

The DTR and DSR signals were originally designed to operate as a handshake mechanism. When a serial port is opened, the DTE device sets its DTR signal to a marked state. Similarly, the DCE sets its DSR signal to a marked state. However, because of the negotiation that takes place with the RTS and CTS signals, the DTR and DSR signals are not commonly used.

The carrier detect and ring indicator signals are used to detect connections with remote modems. These signals are not commonly used.

Signal Polarity

Serial interfaces use a balanced (also called differential) protocol signaling technique. Two serial signals are associated with a circuit: the A signal and the B signal. The A signal is denoted with a plus sign (for example, DTR+), and the B signal is denoted with a minus sign (for example, DTR-). If DTR is low, then DTR+ is negative with respect to DTR-. If DTR is high, then DTR+ is positive with respect to DTR-.

By default, all signal polarities are positive, but sometimes they might be reversed. For example, signals might be miswired as a result of reversed polarities.


Serial Clocking Modes

By default, a serial interface uses loop clocking to determine its timing source. For EIA-530 and V.35 interfaces, you can set each port independently to use one of the following clocking modes. X.21 interfaces can use only loop clocking mode.

- Loop clocking mode—Uses the DCE's receive (RX) clock to clock data from the DCE to the DTE.
- DCE clocking mode—Uses the transmit (TXC) clock, generated by the DCE specifically to be used by the DTE as the DTE's transmit clock.
- Internal clocking mode—Uses an internally generated clock. The speed of this clock is configured locally. Internal clocking mode is also known as line timing.

Both loop clocking mode and DCE clocking mode use external clocks generated by the DCE.

[Figure 36 on page 568](#) shows the clock sources for loop, DCE, and internal clocking modes.

Figure 36: Serial Interface Clocking Modes

Serial Interface Transmit Clock Inversion

When an externally timed clocking mode (DCE or loop) is used, long cables might introduce a phase shift of the DTE-transmitted clock and data. At high speeds, this phase shift might cause errors. Inverting the transmit clock corrects the phase shift, thereby reducing error rates.

DTE Clock Rate Reduction

Although the serial interface is intended for use at the default clock rate of 16.384 MHz, you might need to use a slower rate under any of the following conditions:

- The interconnecting cable is too long for effective operation.
- The interconnecting cable is exposed to an extraneous noise source that might cause an unwanted voltage in excess of +1 volt.

The voltage must be measured differentially between the signal conductor and the point in the circuit from which all voltages are measured ("circuit common") at the load end of the cable, with a 50-ohm resistor substituted for the generator.

- Interference with other signals must be minimized.
- Signals must be inverted.

Serial Line Protocols

Serial interfaces support the following line protocols:

- [EIA-530 on page 569](#)
- [RS-232 on page 569](#)
- [RS-422/449 on page 570](#)
- [V.35 on page 570](#)
- [X.21 on page 571](#)

EIA-530

EIA-530 is an Electronic Industries Association (EIA) standard for the interconnection of DTE and DCE using serial binary data interchange with control information exchanged on separate control circuits. EIA-530 is also known as RS-530.

The EIA-530 line protocol is a specification for a serial interface that uses a DB-25 connector and balanced equivalents of the RS-232 signals—also called V.24. The EIA-530 line protocol is equivalent to the RS-422 and RS-423 interfaces implemented on a 25-pin connector.

The EIA-530 line protocol supports both balanced and unbalanced modes. In unbalanced transmissions, voltages are transmitted over a single wire. Because only a single signal is transmitted, differences in ground potential can cause fluctuations in the measured voltage across the link. For example, if a 3-V signal is sent from one endpoint to another, and the receiving endpoint has a ground potential 1 V higher than the transmitter, the signal on the receiving end is measured as a 2-V signal.

Balanced transmissions use two wires instead of one. Rather than sending a single signal across the wire and having the receiving end measure the voltage, the transmitting device sends two separate signals across two separate wires. The receiving device measures the difference in voltage of the two signals (balanced sampling) and uses that calculation to evaluate the signal. Any differences in ground potential affect both wires equally, and the difference in the signals is still the same.

The EIA-530 interface supports asynchronous and synchronous transmissions at rates ranging from 20 Kbps to 2 Mbps.

RS-232

RS-232 is a Recommended Standard (RS) describing the most widely used type of serial communication. The RS-232 protocol is used for asynchronous data transfer as well as synchronous transfers using HDLC, Frame Relay, and X.25. RS-232 is also known as EIA-232.

The RS-232 line protocol is very popular for low-speed data signals. RS-232 signals are carried as single voltages referred to a common ground signal. The voltage output level of these signals varies between -12 V and $+12\text{ V}$. Within this range, voltages between -3 V and $+3\text{ V}$ are considered inoperative and are used to absorb line noise. Control signals are considered operative when the voltage ranges from $+3\text{ V}$ to $+25\text{ V}$.

The RS-232 line protocol is an unbalanced protocol, because it uses only one wire and is susceptible to signal degradation. Degradation can be extremely disruptive, particularly when a difference in ground potential exists between the transmitting and receiving ends of a link.

The RS-232 interface is implemented in a 25-pin D-shell connector and supports line rates up to 200 Kbps over lines shorter than 98 feet (30 meters).



NOTE: RS-232 serial interfaces cannot function error-free with a clock rate greater than 200 KHz.

RS-422/449

RS-422 is a Recommended Standard (RS) describing the electrical characteristics of balanced voltage digital interface circuits that support higher bandwidths than traditional serial protocols like RS-232. RS-422 is also known as EIA-422.

The RS-449 standard (also known as EIA-449) is compatible with RS-422 signal levels. The EIA created RS-449 to detail the DB-37 connector pinout and define a set of modem control signals for regulating flow control and line status.

The RS-422/499 line protocol runs in balanced mode, allowing serial communications to extend over distances of up to 4,000 feet (1.2 km) and at very fast speeds of up to 10 Mbps.

In an RS-422/499-based system, a single master device can communicate with up to 10 slave devices in the system. To accommodate this configuration, RS-422/499 supports the following kinds of transmission:

- Half-duplex transmission—In half-duplex transmission mode, transmissions occur in only one direction at a time. Each transmission requires a proper handshake before it is sent. This operation is typical of a balanced system in which two devices are connected by a single connection.
- Full-duplex transmission—In full duplex transmission mode, multiple transmissions can occur simultaneously so that devices can transmit and receive at the same time. This operation is essential when a single master in a point-to-multipoint system must communicate with multiple receivers.
- Multipoint transmission—RS-422/449 allows only a single master in a multipoint system. The master can communicate to all points in a multipoint system, and the other points must communicate with each other through the master.

V.35

V.35 is an ITU-T standard describing a synchronous, Physical Layer protocol used for communications between a network access device and a packet network. V.35 is most commonly used in the United States and Europe.

The V.35 line protocol is a mixture of balanced (RS-422) and common ground (RS-232) signal interfaces. The V.35 control signals DTR, DSR, DCD, RTS, and CTS are single-wire common ground signals that are essentially identical to their RS-232 equivalents. Unbalanced signaling for these control signals is sufficient, because the control signals are mostly constant, varying at very low frequency, which makes single-wire transmission suitable. Higher frequency data and clock signals are sent over balanced wires.

V.35 interfaces operate at line rates of 20 Kbps and above.

X.21

X.21 is an ITU-T standard for serial communications over synchronous digital lines. The X.21 protocol is used primarily in Europe and Japan.

The X.21 line protocol is a state-driven protocol that sets up a circuit-switched network using call setup. X.21 interfaces use a 15-pin connector with the following eight signals:

- Signal ground (G)—Reference signal used to evaluate the logic states of the other signals. This signal can be connected to the protective earth (ground).
- DTE common return (Ga)—Reference ground signal for the DCE interface. This signal is used only in unbalanced mode.
- Transmit (T)—Binary signal that carries the data from the DTE to the DCE. This signal can be used for data transfer or in call-control phases such as Call Connect or Call Disconnect.
- Receive (R)—Binary signal that carries the data from the DCE to the DTE. This signal can be used for data transfer or in call-control phases such as Call Connect or Call Disconnect.
- Control (C)—DTE-controlled signal that controls the transmission on an X.21 link. This signal must be on during data transfer, and can be on or off during call-control phases.
- Indication (I)—DCE-controlled signal that controls the transmission on an X.21 link. This signal must be on during data transfer, and can be on or off during call-control phases.
- Signal Element Timing (S)—Clocking signal that is generated by the DCE. This signal specifies when sampling on the line must occur.
- Byte Timing (B)—Binary signal that is on when data or call-control information is being sampled. When an 8-byte transmission is over, this signal switches to off.

Transmissions across an X.21 link require both the DCE and DTE devices to be in a ready state, indicated by an all 1s transmission on the T and R signals.

Related Documentation

- [Example: Configuring a Serial Interface on page 571](#)
- [Example: Deleting a Serial Interface on page 574](#)

Example: Configuring a Serial Interface

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



NOTE: Serial interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

- [Requirements on page 572](#)
- [Overview on page 572](#)

- [Configuration on page 572](#)
- [Verification on page 573](#)

Requirements

Before you begin, install a serial PIM in the SRX Series device. See *SRX Series Services Gateways for the Branch Physical Interface Modules Hardware Guide*.

Overview

In this example, you create the interface se-1/0/0. You create the basic configuration for the new interface by setting the encapsulation type to ppp. Then you set the logical interface to 0. 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. Finally, you set IPv4 address 10.10.10.10/24 on the serial interface.

Configuration

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 se-1/0/0 encapsulation ppp unit 0 family inet address 10.10.10.10/24
```

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 a serial interface:

1. Create the interface.

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

2. Create the basic configuration for the new interface.

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

3. Add logical interfaces.

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

4. Specify an IPv4 address for the interface.

```
[edit interfaces se-1/0/0 unit 0]  
user@host# set family inet address 10.10.10.10/24
```

Results From configuration mode, confirm your configuration by entering the **show interfaces se-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 se-1/0/0
```

```
encapsulation ppp;
unit 0 {
  family inet {
    address 10.10.10.10/24;
  }
}
```

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

Verification

Confirm that the configuration is working properly.

- [Verifying the Link State of All Interfaces on page 573](#)
- [Verifying Interface Properties on page 573](#)

Verifying the Link State of All Interfaces

Purpose Use the ping tool on each peer address in the network to 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 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 se-1/0/0] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the Interfaces> se-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 inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics se-1/0/0** command.

Related Documentation

- [Serial Interfaces Overview on page 565](#)
- [Example: Deleting a Serial Interface on page 574](#)

Example: Deleting a Serial Interface

This example shows how to delete a serial interface.



NOTE: Serial interfaces are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Requirements

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

Overview

In this example, you delete the se-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 J-Web pages.

Configuration

Step-by-Step Procedure

To delete a serial interface:

1. Specify the interface you want to delete.

```
[edit]
user@host# delete se-1/0/0
```
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.

Related Documentation

- [Serial Interfaces Overview on page 565](#)
- [Example: Configuring a Serial Interface on page 571](#)

Understanding the 8-Port Synchronous Serial GPIM

A Gigabit-Backplane Physical Interface Module (GPIM) 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: Serial interfaces, including the 8-port synchronous serial GPIM, are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

The 8-port synchronous serial GPIM provides the physical connection to serial network media types, receiving incoming packets from the network and transmitting outgoing packets to the network. Besides forwarding packets for processing, the GPIM performs framing and line-speed signaling. This GPIM provides 8 ports that operate in sync mode and supports a line rate of 64 Mbps or 8 Mbps per port.

Supported Features

Table 49 on page 575 lists the features supported on the 8-port synchronous serial GPIM.

Table 49: Supported Features

Features	Description
Operation modes (autoselection based on cable, no configuration required)	<ul style="list-style-type: none"> • DTE (data terminal equipment) • DCE (data communication equipment)

Table 49: Supported Features (continued)

Features	Description
Clocking	<ul style="list-style-type: none"> Tx clock modes <ul style="list-style-type: none"> DCE clock (only valid in DTE mode) Baud clock (internally generated) Loop clock (external) Rx clock modes <ul style="list-style-type: none"> Baud clock (internally generated) Loop clock (external)
Clock rates (baud rates)	1.2 KHz to 8.0 MHz NOTE: RS-232 serial interfaces might cause an error with a clock rate greater than 200 KHz.
MTU	9192 bytes, default value is 1504 bytes
HDLC features	<ul style="list-style-type: none"> Idle flag/fill (0x7e or all ones), default idle flag is (0x7e) Counters—giants, runts, FCS error, abort error, align error
Line encoding	NRZ and NRZI
Invert data	Enabled
Line protocol	EIA530/EIA530A, X.21, RS-449, RS-232, V.35
Data cables	Separate cable for each line protocol (both DTE/DCE mode)
Error counters (conformance to ANSI specification)	Enabled
Alarms and defects	<ul style="list-style-type: none"> Rx clock absent Tx clock absent DCD absent RTS/CTS absent DSR/DTR absent
Data signal	Rx clock
Control signals	<ul style="list-style-type: none"> To DTE: CTS, DCD, DSR From DTE: DTR, RTS
Serial autoresync	<ul style="list-style-type: none"> Configurable resync duration Configurable resync interval
Diagnostic features	<ul style="list-style-type: none"> Loopback modes—local, remote, and dce-local loopback Ability to ignore control signals

Table 49: Supported Features (continued)

Features	Description
Layer 2 features	Encapsulation <ul style="list-style-type: none"> • PPP • Cisco HDLC • Frame Relay • MLPPP • MLFR
SNMP features	SNMP information receivable at each port <ul style="list-style-type: none"> • IF-MIB - rfc2863a.mib • jnx-chassis.mib
Anticounterfeit check	Enabled

**Related
Documentation**

- [Example: Configuring an 8-Port Synchronous Serial GPIM in Back-to-Back SRX650 Services Gateways on page 577](#)

Example: Configuring an 8-Port Synchronous Serial GPIM in Back-to-Back SRX650 Services Gateways

This example shows how to perform a basic back-to-back device configuration with an 8-port synchronous serial GPIM. It describes the most common scenario in which a serial GPIM is deployed.

In this example, the SRX650 devices are shown as both data communication equipment (DCE) and data terminal equipment (DTE). In certain deployment scenarios, the DTE can be a serial modem or an encryptor or decryptor.



NOTE: Serial interfaces, including the 8-port synchronous serial GPIM, are no longer supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

- [Requirements on page 578](#)
- [Overview and Topology on page 578](#)
- [Configuration on page 579](#)
- [Verification on page 587](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.1 R2 or later for SRX Series Services Gateways.
- Two SRX650 devices connected back-to-back.
- Two 8-port synchronous serial GPIMs.
- Four pairs of DCE and DTE cables. The cable can be any type as mentioned in *8-Port Serial GPIM Interface Cables*.

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 265](#).

Overview and Topology


In this scenario, the configuration is done on two interfaces. All ports are configured with different encapsulations, such as Cisco High-Level Data Link Control (HDLC), Frame Relay, and Point-to-Point Protocol (PPP). When Frame Relay is set, then the data link connection identifier (in this example, 111) must also be set.

In this example, all eight ports on Device 1 (SRX650) are configured in DTE mode and their respective eight ports on Device 2 (SRX650) are configured in DCE mode.

For Device 1, you set the encapsulation type to **ppp**. Then you set the logical interface to **0**. 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. Finally, you set the IPv4 address to 10.10.10.1/24 on the serial port. For Device 2, you follow a procedure similar to Device 1, but you set the clocking mode to dce.

[Figure 37 on page 579](#) shows the topology used in this example.

Figure 37: Basic Back-to-Back Device Configuration



g034406

Configuration

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

```
Device 1 set interfaces se-7/0/0 mtu 9192
set interfaces se-7/0/0 encapsulation ppp
set interfaces se-7/0/0 serial-options clocking-mode internal
set interfaces se-7/0/0 unit 0 family inet address 10.10.10.1/24
set interfaces se-7/0/1 mtu 9192
set interfaces se-7/0/1 encapsulation cisco-hdlc
set interfaces se-7/0/1 serial-options clocking-mode internal
set interfaces se-7/0/1 unit 0 family inet address 11.11.11.1/24
set interfaces se-7/0/2 dce
set interfaces se-7/0/2 mtu 9192
set interfaces se-7/0/2 encapsulation frame-relay
set interfaces se-7/0/2 serial-options clocking-mode internal
set interfaces se-7/0/2 unit 0 dlci 111
set interfaces se-7/0/2 unit 0 family inet address 12.12.12.1/24
set interfaces se-7/0/3 mtu 9192
set interfaces se-7/0/3 encapsulation ppp
set interfaces se-7/0/3 serial-options clocking-mode internal
set interfaces se-7/0/3 unit 0 family inet address 13.13.13.1/24
```

```

set interfaces se-7/0/4 mtu 9192
set interfaces se-7/0/4 encapsulation cisco-hdlc
set interfaces se-7/0/4 serial-options clocking-mode internal
set interfaces se-7/0/4 unit 0 family inet address 14.14.14.1/24
set interfaces se-7/0/5 dce
set interfaces se-7/0/5 mtu 9192
set interfaces se-7/0/5 encapsulation frame-relay
set interfaces se-7/0/5 serial-options clocking-mode internal
set interfaces se-7/0/5 unit 0 dlci 112
set interfaces se-7/0/5 unit 0 family inet address 15.15.15.1/24
set interfaces se-7/0/6 mtu 9192
set interfaces se-7/0/6 encapsulation cisco-hdlc
set interfaces se-7/0/6 serial-options clocking-mode internal
set interfaces se-7/0/6 unit 0 family inet address 16.16.16.1/24
set interfaces se-7/0/7 mtu 9192
set interfaces se-7/0/7 encapsulation ppp
set interfaces se-7/0/7 serial-options clocking-mode internal
set interfaces se-7/0/7 unit 0 family inet address 17.17.17.1/24
set routing-options static route 21.21.21.0/24 next-hop 10.10.10.2
set routing-options static route 23.23.23.0/24 next-hop 11.11.11.2
set routing-options static route 25.25.25.0/24 next-hop 12.12.12.2
set routing-options static route 27.27.27.0/24 next-hop 13.13.13.2
set routing-options static route 29.29.29.0/24 next-hop 14.14.14.2
set routing-options static route 31.31.31.0/24 next-hop 15.15.15.2
set routing-options static route 33.33.33.0/24 next-hop 16.16.16.2
set routing-options static route 35.35.35.0/24 next-hop 17.17.17.2

```

Device 2

```

set interfaces se-3/0/0 mtu 9192
set interfaces se-3/0/0 encapsulation ppp
set interfaces se-3/0/0 serial-options clocking-mode dce
set interfaces se-3/0/0 unit 0 family inet address 10.10.10.2/24
set interfaces se-3/0/1 mtu 9192
set interfaces se-3/0/1 encapsulation cisco-hdlc
set interfaces se-3/0/1 serial-options clocking-mode dce
set interfaces se-3/0/1 unit 0 family inet address 11.11.11.2/24
set interfaces se-3/0/2 dce
set interfaces se-3/0/2 mtu 9192
set interfaces se-3/0/2 encapsulation frame-relay
set interfaces se-3/0/2 serial-options clocking-mode dce
set interfaces se-3/0/2 unit 0 dlci 111
set interfaces se-3/0/2 unit 0 family inet address 12.12.12.2/24
set interfaces se-3/0/3 mtu 9192
set interfaces se-3/0/3 encapsulation ppp
set interfaces se-3/0/3 serial-options clocking-mode dce
set interfaces se-3/0/3 unit 0 family inet address 13.13.13.2/24
set interfaces se-3/0/4 mtu 9192
set interfaces se-3/0/4 encapsulation cisco-hdlc
set interfaces se-3/0/4 serial-options clocking-mode dce
set interfaces se-3/0/4 unit 0 family inet address 14.14.14.2/24
set interfaces se-3/0/5 dce
set interfaces se-3/0/5 mtu 9192
set interfaces se-3/0/5 encapsulation frame-relay
set interfaces se-3/0/5 serial-options clocking-mode dce
set interfaces se-3/0/5 unit 0 dlci 112
set interfaces se-3/0/5 unit 0 family inet address 15.15.15.2/24

```

```

set interfaces se-3/0/6 mtu 9192
set interfaces se-3/0/6 encapsulation cisco-hdlc
set interfaces se-3/0/6 serial-options clocking-mode dce
set interfaces se-3/0/6 unit 0 family inet address 16.16.16.2/24
set interfaces se-3/0/7 mtu 9192
set interfaces se-3/0/7 encapsulation ppp
set interfaces se-3/0/7 serial-options clocking-mode dce
set interfaces se-3/0/7 unit 0 family inet address 17.17.17.2/24
set routing-options static route 20.20.20.0/24 next-hop 10.10.10.1
set routing-options static route 22.22.22.0/24 next-hop 11.11.11.1
set routing-options static route 24.24.24.0/24 next-hop 12.12.12.1
set routing-options static route 26.26.26.0/24 next-hop 13.13.13.1
set routing-options static route 28.28.28.0/24 next-hop 14.14.14.1
set routing-options static route 30.30.30.0/24 next-hop 15.15.15.1
set routing-options static route 32.32.32.0/24 next-hop 16.16.16.1
set routing-options static route 34.34.34.0/24 next-hop 17.17.17.1

```

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 the interfaces on Device 1:

1. Specify the maximum transmission unit (MTU) value for the interface.

```

[edit interfaces]
user@host# set se-7/0/0 mtu 9192

```

2. Configure the encapsulation type.

```

[edit interfaces]
user@host# set se-7/0/0 encapsulation ppp

```

3. Configure the serial options, such as the clocking mode.

```

[edit interfaces]
user@host# set se-7/0/0 serial-options clocking-mode internal

```

4. Set the IPv4 address on the serial port.

```

[edit interfaces]
user@host# set se-7/0/0 unit 0 family inet address 10.10.10.1/24

```

5. Configure the static route information.

```

[edit routing-options]
user@host# set static route 21.21.21.0/24 next-hop 10.10.10.2

```



NOTE: Repeat the same configuration for the other seven ports on Device 1.

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

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

Step-by-Step Procedure

To configure the interfaces on Device 2:

1. Specify the MTU value for the interface.

```
[edit interfaces]  
user@host# set se-3/0/0 mtu 9192
```

2. Configure the encapsulation type.

```
[edit interfaces]  
user@host# set se-3/0/0 encapsulation ppp
```

3. Configure the serial options, such as the clocking mode.

```
[edit interfaces]  
user@host# set se-3/0/0 serial-options clocking-mode dce
```

4. Set the IPv4 address on the serial port.

```
[edit interfaces]  
user@host# set se-3/0/0 unit 0 family inet address 10.10.10.2/24
```

5. Configure the static route information.

```
[edit routing-options]  
user@host# set static route 20.20.20.0/24 next-hop 10.10.10.1
```



NOTE: Repeat the same configuration for the other seven ports on Device 2.

6. 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.

```
Device 1 [edit]  
user@host# show interfaces
```

```
se-7/0/0 {
  mtu 9192;
  encapsulation ppp;
  serial-options {
    clocking-mode internal;
  }
  unit 0 {
    family inet {
      address 10.10.10.1/24;
    }
  }
}
se-7/0/1 {
  mtu 9192;
  encapsulation cisco-hdlc;
  serial-options {
    clocking-mode internal;
  }
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
  }
}
se-7/0/2 {
  dce;
  mtu 9192;
  encapsulation frame-relay;
  serial-options {
    clocking-mode internal;
  }
  unit 0 {
    dlci 111;
    family inet {
      address 12.12.12.1/24;
    }
  }
}
se-7/0/3 {
  mtu 9192;
  encapsulation ppp;
  serial-options {
    clocking-mode internal;
  }
  unit 0 {
    family inet {
      address 13.13.13.1/24;
    }
  }
}
se-7/0/4 {
  mtu 9192;
  encapsulation cisco-hdlc;
  serial-options {
    clocking-mode internal;
  }
}
```

```
    unit 0 {
        family inet {
            address 14.14.14.1/24;
        }
    }
}
se-7/0/5 {
    dce;
    mtu 9192;
    encapsulation frame-relay;
    serial-options {
        clocking-mode internal;
    }
    unit 0 {
        dlci 112;
        family inet {
            address 15.15.15.1/24;
        }
    }
}
se-7/0/6 {
    mtu 9192;
    encapsulation cisco-hdlc;
    serial-options {
        clocking-mode internal;
    }
    unit 0 {
        family inet {
            address 16.16.16.1/24;
        }
    }
}
se-7/0/7 {
    mtu 9192;
    encapsulation ppp;
    serial-options {
        clocking-mode internal;
    }
    unit 0 {
        family inet {
            address 17.17.17.1/24;
        }
    }
}
```

[edit]

user@host# show routing-options

```
static {
    route 21.21.21.0/24 next-hop 10.10.10.2;
    route 23.23.23.0/24 next-hop 11.11.11.2;
    route 25.25.25.0/24 next-hop 12.12.12.2;
    route 27.27.27.0/24 next-hop 13.13.13.2;
    route 29.29.29.0/24 next-hop 14.14.14.2;
    route 31.31.31.0/24 next-hop 15.15.15.2;
    route 33.33.33.0/24 next-hop 16.16.16.2;
    route 35.35.35.0/24 next-hop 17.17.17.2;
```



```
}

```

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

```
Device 2 [edit]
user@host# show interfaces
se-3/0/0 {
  mtu 9192;
  encapsulation ppp;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    family inet {
      address 10.10.10.2/24;
    }
  }
}
se-3/0/1 {
  mtu 9192;
  encapsulation cisco-hdlc;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    family inet {
      address 11.11.11.2/24;
    }
  }
}
se-3/0/2 {
  dce;
  mtu 9192;
  encapsulation frame-relay;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    dlci 111;
    family inet {
      address 12.12.12.2/24;
    }
  }
}
se-3/0/3 {
  mtu 9192;
  encapsulation ppp;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    family inet {
      address 13.13.13.2/24;
    }
  }
}

```

```
se-3/0/4 {
  mtu 9192;
  encapsulation cisco-hdlc;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    family inet {
      address 14.14.14.2/24;
    }
  }
}
se-3/0/5 {
  dce;
  mtu 9192;
  encapsulation frame-relay;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    dlci 112;
    family inet {
      address 15.15.15.2/24;
    }
  }
}
se-3/0/6 {
  mtu 9192;
  encapsulation cisco-hdlc;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    family inet {
      address 16.16.16.2/24;
    }
  }
}
se-3/0/7 {
  mtu 9192;
  encapsulation ppp;
  serial-options {
    clocking-mode dce;
  }
  unit 0 {
    family inet {
      address 17.17.17.2/24;
    }
  }
}
```

[edit]

user@host# show routing-options

```
static {
  route 20.20.20.0/24 next-hop 10.10.10.1;
  route 22.22.22.0/24 next-hop 11.11.11.1;
```

```

route 24.24.24.0/24 next-hop 12.12.12.1;
route 26.26.26.0/24 next-hop 13.13.13.1;
route 28.28.28.0/24 next-hop 14.14.14.1;
route 30.30.30.0/24 next-hop 15.15.15.1;
route 32.32.32.0/24 next-hop 16.16.16.1;
route 34.34.34.0/24 next-hop 17.17.17.1;
}

```

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

Verification

Confirm that the configuration is working properly.

- [Verifying Interface Link Status on page 587](#)
- [Verifying Interface Statistics for DCE on page 587](#)
- [Verifying Interface Statistics for DTE on page 590](#)

Verifying Interface Link Status

Purpose Verify that the interface link status is up.

Action From operational mode, enter the **show interface terse se-7/0/*** command.

```
user@srx650-1> show interface terse se-7/0/*
```

Interface	Admin	Link	Proto	Local	Remote
se-7/0/0	up	up			
se-7/0/0.0	up	up	inet	10.10.10.1/24	
se-7/0/1	up	up			
se-7/0/1.0	up	up	inet	11.11.11.1/24	
se-7/0/2	up	up			
se-7/0/2.0	up	up	inet	12.12.12.1/24	
se-7/0/3	up	up			
se-7/0/3.0	up	up	inet	13.13.13.1/24	
se-7/0/4	up	up			
se-7/0/4.0	up	up	inet	14.14.14.1/24	
se-7/0/5	up	up			
se-7/0/5.0	up	up	inet	15.15.15.1/24	
se-7/0/6	up	up			
se-7/0/6.0	up	up	inet	16.16.16.1/24	
se-7/0/7	up	up			
se-7/0/7.0	up	up	inet	17.17.17.1/24	

Meaning The output displays a list of all interfaces configured. If the Link column displays **up** for all interfaces, the configuration is working properly. This verifies that the GPIM is up and end-to-end ping is working.

Verifying Interface Statistics for DCE

Purpose Verify that the interfaces are configured properly for DCE.

Action From operational mode, enter the **show interface se-7/0/0 extensive | no-more** command.

```
user@srx650-1>show interface se-7/0/0 extensive | no-more
```

```
Physical interface: se-7/0/0, Enabled, Physical link is Up
Interface index: 161, SNMP ifIndex: 592, Generation: 164
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 8mbps
Device flags : Present Running
Interface flags: Point-To-Point Internal: 0x0
Link flags : Keepalives
Hold-times : Up 0 ms, Down 0 ms
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 123 (last seen 00:00:02 ago)
  Output: 123 (last sent 00:00:01 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
CoS queues : 8 supported, 8 maximum usable queues
Last flapped : 2011-06-27 22:57:24 PDT (00:20:59 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :                23792                160 bps
  Output bytes :                22992                536 bps
  Input packets:                 404                 0 pps
  Output packets:                409                 0 pps
Input errors:
  Errors: 3, Drops: 0, Framing errors: 3, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
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            409            409              0

Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Serial media information:
  Line protocol: eia530
  Resync history:
    Sync loss count: 0
  Data signal:
    Rx Clock: OK
  Control signals:
  Local mode: DCE
    To DTE: CTS: up, DCD: up, DSR: up
    From DTE: DTR: up, RTS: up
  DCE loopback override: Off
```

```

Clocking mode: internal
Loopback: none
Tx clock: non-invert
Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
    0 best-effort         95      7600000    95      0      low
none
    3 network-control     5       400000    5       0      low
none

Logical interface se-7/0/0.0 (Index 82) (SNMP ifIndex 600) (Generation 147)
Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPP
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp
Flow Statistics :
Flow Input statistics :
  Self packets :          153
  ICMP packets :           0
  VPN packets :           0
  Multicast packets :      0
  Bytes permitted by policy : 13152
  Connections established : 1
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: 162, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
  Generation: 175

```

Meaning The output displays a list of all DCE verification parameters and the mode configured. If the local mode displays DCE, the configuration is working properly.

Verifying Interface Statistics for DTE

Purpose Verify that the interfaces are configured properly for DTE.

Action From operational mode, enter the `show interfaces se-3/0/0 extensive | no-more` command.

```
user@srx650-2>show interfaces se-3/0/0 extensive | no-more
```

```
Physical interface: se-3/0/0, Enabled, Physical link is Up
Interface index: 168, SNMP ifIndex: 594, Generation: 171
Type: Serial, Link-level type: PPP, MTU: 1504, Maximum speed: 8mbps
Device flags   : Present Running
Interface flags: Point-To-Point Internal: 0x0
Link flags     : Keepalives
Hold-times    : Up 0 ms, Down 0 ms
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
  Input : 242 (last seen 00:00:09 ago)
  Output: 242 (last sent 00:00:10 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Closed
PAP state: Closed
CoS queues   : 8 supported, 8 maximum usable queues
Last flapped : 2011-06-27 22:52:06 PDT (00:40:41 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :          44582          0 bps
  Output bytes :          42872          0 bps
  Input packets:           776          0 pps
  Output packets:          779          0 pps
Input errors:
  Errors: 6, Drops: 0, Framing errors: 6, Runts: 0, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              2              2              0
  1 expedited-fo             0              0              0
  2 assured-forw             0              0              0
  3 network-cont            777            777              0

Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Serial media information:
```

```

Line protocol: eia530
Resync history:
  Sync loss count: 0
Data signal:
  Rx Clock: OK
Control signals:
  Local mode: DTE
  To DCE: DTR: up, RTS: up
  From DCE: CTS: up, DCD: up, DSR: up
Clocking mode: loop-timed
Loopback: none
Tx clock: non-invert
Line encoding: nrz
Packet Forwarding Engine configuration:
  Destination slot: 3
CoS information:
  Direction : Output
  CoS transmit queue
Limit
                                %          bps          %          usec          Priority
0 best-effort                  95          7600000      95           0           low
none
3 network-control              5           400000       5            0           low
none
Logical interface se-3/0/0.0 (Index 82) (SNMP ifIndex 602) (Generation 147)
Flags: Point-To-Point SNMP-Traps 0x0 Encapsulation: PPP
Security: Zone: HOST
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf pgm pim rip router-discovery rsvp sap vrrp
Flow Statistics :
Flow Input statistics :
  Self packets :                287
  ICMP packets :                0
  VPN packets :                 0
  Multicast packets :           0
  Bytes permitted by policy :    24044
  Connections established :      1
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: 162, Route table: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.2, Broadcast: 10.10.10.255,
Generation: 175
    
```

Meaning The output displays a list of all DTE verification parameters and the mode configured. If the local mode displays DTE, the configuration is working properly.

Related Documentation

- [Understanding the 8-Port Synchronous Serial GPIM on page 575](#)

PART 9

Configuration Statements and Operational Commands

- [Configuration Statements on page 595](#)
- [Operational Commands on page 683](#)

CHAPTER 38

Configuration Statements

- [accept-source-mac](#) on page 598
- [access-point-name](#) on page 599
- [apply-groups](#) on page 599
- [arp-resp](#) on page 600
- [authentication-method \(Interfaces\)](#) on page 600
- [bandwidth \(Interfaces\)](#) on page 601
- [bundle \(Interfaces\)](#) on page 601
- [cbr rate](#) on page 602
- [cellular-options](#) on page 603
- [classifiers \(CoS\)](#) on page 604
- [client-identifier \(Interfaces\)](#) on page 605
- [code-points \(CoS\)](#) on page 606
- [compression-device \(Interfaces\)](#) on page 607
- [credit \(Interfaces\)](#) on page 607
- [data-rate](#) on page 608
- [disable \(PoE\)](#) on page 609
- [dhcp \(Interfaces\)](#) on page 610
- [duration \(PoE\)](#) on page 611
- [family inet \(Interfaces\)](#) on page 612
- [family inet6](#) on page 615
- [flag \(Interfaces\)](#) on page 618
- [flexible-vlan-tagging \(Interfaces\)](#) on page 619
- [flow-control \(Interfaces\)](#) on page 619
- [flow-monitoring \(Services\)](#) on page 620
- [forwarding-classes \(CoS\)](#) on page 621
- [fpc \(Interfaces\)](#) on page 622
- [gratuitous-arp-reply](#) on page 623
- [gsm-options](#) on page 624

- [guard-band \(PoE\) on page 625](#)
- [hub-assist on page 625](#)
- [inline-jflow \(Forwarding Options\) on page 626](#)
- [interface \(PIC Bundle\) on page 626](#)
- [interface \(PoE\) on page 627](#)
- [interfaces \(CoS\) on page 628](#)
- [interval \(Interfaces\) on page 629](#)
- [interval \(PoE\) on page 629](#)
- [ipv4-template \(Services\) on page 630](#)
- [ipv6-template \(Services\) on page 630](#)
- [lACP \(Interfaces\) on page 631](#)
- [latency \(Interfaces\) on page 631](#)
- [lease-time on page 632](#)
- [line-rate \(Interfaces\) on page 633](#)
- [link-speed \(Interfaces\) on page 633](#)
- [loopback \(Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet\) on page 634](#)
- [loss-priority \(CoS Loss Priority\) on page 635](#)
- [loss-priority \(CoS Rewrite Rules\) on page 636](#)
- [loss-priority-maps \(CoS Interfaces\) on page 637](#)
- [loss-priority-maps \(CoS\) on page 637](#)
- [management \(PoE\) on page 638](#)
- [maximum-power \(PoE\) on page 639](#)
- [media-type \(Interfaces\) on page 639](#)
- [minimum-links \(Interfaces\) on page 640](#)
- [mtu on page 641](#)
- [native-vlan-id on page 642](#)
- [next-hop-tunnel on page 644](#)
- [no-dns-propagation on page 645](#)
- [option-refresh-rate \(Services\) on page 645](#)
- [pic-mode \(Chassis T1 Mode\) on page 646](#)
- [periodic \(Interfaces\) on page 647](#)
- [ppp-over-ether on page 648](#)
- [pppoe on page 649](#)
- [pppoe-options on page 650](#)
- [priority \(PoE\) on page 651](#)
- [profile \(Access\) on page 652](#)
- [profiles on page 654](#)

- [promiscuous-mode \(Interfaces\) on page 655](#)
- [quality \(Interfaces\) on page 655](#)
- [r2cp on page 656](#)
- [radio-router \(Interfaces\) on page 657](#)
- [redundancy-group \(Interfaces\) on page 658](#)
- [redundant-ether-options on page 659](#)
- [redundant-parent \(Interfaces Fast Ethernet\) on page 660](#)
- [redundant-parent \(Interfaces Gigabit Ethernet\) on page 661](#)
- [request pppoe connect](#)
- [request pppoe disconnect](#)
- [resource \(Interfaces\) on page 664](#)
- [retransmission-attempt on page 664](#)
- [retransmission-interval \(Interfaces\) on page 665](#)
- [roaming-mode on page 665](#)
- [scheduler-map \(CoS Virtual Channels\) on page 666](#)
- [select-profile on page 666](#)
- [server-address on page 667](#)
- [shaping-rate \(CoS Interfaces\) on page 668](#)
- [simple-filter \(Interfaces\) on page 669](#)
- [sip-password on page 669](#)
- [sip-user-id on page 670](#)
- [source-address-filter \(Interfaces\) on page 671](#)
- [source-filtering \(Interfaces\) on page 672](#)
- [speed \(Interfaces\) on page 672](#)
- [speed \(Gigabit Ethernet interface\) on page 673](#)
- [speed \(Chassis Cluster\) on page 675](#)
- [telemetries \(PoE\) on page 676](#)
- [template-refresh-rate \(Services\) on page 677](#)
- [threshold \(Interfaces\) on page 677](#)
- [traceoptions \(Interfaces\) on page 678](#)
- [update-server on page 678](#)
- [vbr rate on page 679](#)
- [vdsl-profile on page 680](#)
- [vendor-id \(Interfaces\) on page 681](#)
- [web-authentication \(Interfaces\) on page 682](#)

accept-source-mac

Syntax `accept-source-mac {
 mac-address mac-address;
 }`

Hierarchy Level [edit interfaces *interface-name* unit logical-unit-number]

Release Information Statement introduced in Junos OS Release 11.4.

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.

Related Documentation • [Understanding Ethernet Interfaces on page 259](#)

access-point-name

Syntax	<code>access-point-name <i>apn</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
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	<i>apn</i> —Access point name.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

apply-groups

Syntax	<code>apply-groups;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router]
Release Information	Statement introduced in Junos OS Release 9.6. Statement modified in Junos OS Release 15.1.
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.
Related Documentation	<ul style="list-style-type: none"> • Configuring PPPoE-Based Radio-to-Router Protocols on page 418

arp-resp

Syntax	arp-resp (restricted unrestricted);
Hierarchy Level	[edit interfaces <i>interfaces-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Configure Address Resolution Protocol (ARP) response on the interface.
Options	<ul style="list-style-type: none">• restricted—Enable restricted proxy ARP response on the interface. This is the default.• unrestricted—Enable unrestricted ARP response on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Interfaces Feature Guide for Security Devices</i>

authentication-method (Interfaces)

Syntax	authentication-method (pap chap none);
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Specify the authentication method for connection to a Global System for Mobile Communications (GSM) cellular network.
Options	<ul style="list-style-type: none">• 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.

bandwidth (Interfaces)

Syntax	<code>bandwidth <i>bandwidth</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> radio-router]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This option controls the weight of the current (vs. maximum) data rate (value 0–100).
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 415

bundle (Interfaces)

Syntax	<code>bundle <i>bundle-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mlppp]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Specify the logical interface name the link joins.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

cbr rate

Syntax	cbr rate;
Hierarchy Level	[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
Release Information	Command introduced in Release 9.5 of Junos OS.
Description	For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.
Options	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Interfaces Configuration Guide for Security Devices</i>

cellular-options

```
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*]

Release Information Statement introduced in Junos OS Release 9.5.

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 interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.


classifiers (CoS)

Syntax	<pre> classifiers { (dscp dscp-ipv6 exp ieee-802.1 ieee-802.1ad inet-precedence) <i>classifier-name</i> { forwarding-class <i>forwarding-class-name</i> { loss-priority (high low medium-high medium-low) { code-point <i>alias-or-bit-string</i> ; } import (default <i>user-defined</i>); } } } </pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in Junos OS Release 9.2
Description	Configure a user-defined behavior aggregate (BA) classifier.
Options	<ul style="list-style-type: none"> <i>classifier-name</i>—User-defined name for the classifier. import (default <i>user-defined</i>)—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 <i>class-name</i>—Specify the name of the forwarding class. You can use the default forwarding class names or define new ones. loss-priority <i>level</i>—Specify a loss priority for this forwarding class: high, low, medium-high, medium-low. code-points (<i>alias</i> <i>bits</i>)—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.
Related Documentation	<ul style="list-style-type: none"> Understanding Interfaces on page 3

client-identifier (Interfaces)

Syntax	<pre>client-identifier { (ascii <i>string</i> hexadecimal <i>string</i>); }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i> dhcp]
Release Information	Statement introduced in Junos OS Release 9.2.
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	<ul style="list-style-type: none">• ascii <i>ascii</i> —Identifier consisting of ASCII characters.• hexadecimal <i>hexadecimal</i> —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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

code-points (CoS)

Syntax	<code>code-points ([<i>aliases</i>] [<i>bit-patterns</i>]);</code>
Hierarchy Level	[edit class-of-service classifiers <i>type classifier-name</i> forwarding-class <i>class-name</i> loss-priority <i>level</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.1X44 for the SRX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 14.2 for PTX Series Packet Transport Routers.
Description	Specify one or more DSCP code-point aliases or bit sets to apply to a forwarding class.. <div> NOTE: OCX Series switches do not support MPLS, and therefore, do not support EXP code points or code point aliases.</div>
Options	<i>aliases</i> —Name of the DSCP alias. <i>bit-patterns</i> —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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3• <i>Understanding How Behavior Aggregate Classifiers Prioritize Trusted Traffic</i>• <i>Example: Configuring Behavior Aggregate Classifiers</i>• <i>Example: Configuring BA Classifiers on Transparent Mode Security Devices</i>

compression-device (Interfaces)

Syntax	<code>compression-device <i>name</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit (Interfaces) <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify the compression interface for voice services traffic.
Options	<i>name</i> —Name of the AC.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

credit (Interfaces)

Syntax	<pre>credit { interval <i>number</i>; }</pre>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> radio—router]</code>
Release Information	Statement introduced in Junos OS Release 10.1.
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.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

data-rate

Syntax	<code>data-rate <i>weight</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router]
Release Information	Statement introduced in Release 10.2 of Junos OS .
Description	Configure the weight of the resource factor when calculating an effective data rate.
Options	<i>weight</i> —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.
Related Documentation	<ul style="list-style-type: none">• Configuring PPPoE-Based Radio-to-Router Protocols on page 418

disable (PoE)

Syntax	disable;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)] [edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	<p>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.</p>
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.
Related Documentation	<ul style="list-style-type: none">• Example: Disabling a PoE Interface on page 376

dhcp (Interfaces)

Syntax	<pre>dhcp { client-identifier { (ascii <i>string</i> hexadecimal <i>string</i>); } lease-time (<i>length</i> infinite); retransmission-attempt <i>value</i>; retransmission-interval <i>seconds</i>; server-address <i>server-address</i>; update-server; vendor-id <i>vendor-id</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Configure the Dynamic Host Configuration Protocol (DHCP) client.
Options	The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

duration (PoE)

Syntax	<code>duration <i>hours</i>;</code>
Hierarchy Level	[edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in Junos OS Release 9.5.
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	<p>hours— Hours for which telemetry data should be retained.</p> <p>Range: 1 through 24 hours</p> <p>Default: 1 hour</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring PoE on All Interfaces on page 370

family inet (Interfaces)

```
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*]

Release Information Statement supported in Junos 10.2 for SRX Series devices.

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.

Related Documentation

- [Understanding Interfaces on page 3](#)

family inet6

```
Syntax  inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
    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*]

Release Information Statement supported in Junos 10.2 for SRX Series devices.

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.

Related Documentation • [Understanding Interfaces on page 3](#)

flag (Interfaces)

Syntax	flag
Hierarchy Level	[edit interfaces interface-name traceoptions]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Define tracing operations for individual interfaces. To specify more than one tracing operation, include multiple flag statements.
Options	<ul style="list-style-type: none">• 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.
	<div> NOTE:<ul style="list-style-type: none">• MTU is limited to 1518 on this interface.• Cache and enhanced options are applicable only to Enhanced Web Filtering.</div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

flexible-vlan-tagging (Interfaces)

Syntax	<code>flexible-vlan-tagging;</code>
Hierarchy Level	[edit interfaces <i>interface</i>]
Release Information	Statement introduced in Junos OS Release 12.1X44-D10.
Description	Simultaneously supports transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.
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.
Related Documentation	<ul style="list-style-type: none"> • Configuring VLAN Tagging on page 57

flow-control (Interfaces)

Syntax	<code>(flow-control no-flow-control);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> fastether-options] [edit interfaces <i>interface-name</i> gigether-options] [edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement modified in Junos OS Release 9.2.
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.
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 259

flow-monitoring (Services)

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]

Release Information Statement introduced in Junos OS Release 10.4.

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.

Related Documentation

- [Understanding Interfaces on page 3](#)

forwarding-classes (CoS)

Syntax

```
forwarding-classes {
  class class-name {
    priority (high | low);
    queue-num number;
    spu-priority (high | low | medium-high | medium-low);
  }
  queue queue-number {
    class-name {
      priority (high | low);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced in Junos OS Release 8.5. Statement updated in Junos OS Release 11.4. The **spu-priority** option introduced in Junos OS Release 11.4R2.

Description Configure forwarding classes and assign queue numbers.

Options

- **class *class-name***—Display the forwarding class name assigned to the internal queue number.



NOTE: This option is supported only on SRX1500, SRX5400, SRX5600, and SRX5800.



NOTE: AppQoS forwarding classes must be different from those defined for interface-based rewriters.

- **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**.

- **queue *queue-number***—Specify the internal queue number to which a forwarding class is assigned.
- **spu-priority**—Services Processing Unit (SPU) priority queue, **high**, **medium-high**, **medium-low**, or **low**. The default **spu-priority** is **low**.



NOTE: The `spu-priority` option is only supported on SRX1500 devices and SRX5000 line devices.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring AppQoS*

fpc (Interfaces)

Syntax `fpc;`

Hierarchy Level [edit interfaces pic-set *pic-set-name*]

Release Information Command introduced in Junos OS Release 9.6.

Description Sets the PIC bundle and the FPC slot.

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.

Related Documentation

- [Understanding Interfaces on page 3](#)

gratuitous-arp-reply

Syntax	(gratuitous-arp-reply no-gratuitous-arp-reply);
Hierarchy Level	[edit interfaces <i>interface-name</i>] [edit interfaces interface-range <i>interface-range-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 in EX Series switches. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
Description	For 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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Gratuitous ARP</i>• <i>no-gratuitous-arp-request</i>

gsm-options

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]`

Release Information Statement introduced in Junos OS Release 9.5.

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 interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

guard-band (PoE)

Syntax	<code>guard-band <i>watts</i>;</code>
Hierarchy Level	[edit poe]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Reserves the specified amount of power for the SRX Series device in case of a spike in PoE consumption.
Options	<p>watts—Amount of power to be reserved for the SRX Series device in case of a spike in PoE consumption.</p> <p>Range: 0 through 19 W</p> <p>Default: 0 W</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Power over Ethernet on page 367

hub-assist

Syntax	<code>hub-assist <i>weight</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Configure the weight of the resource factor when calculating an effective interface bandwidth.
Options	<p>weight—Factor used to calculate interface bandwidth.</p> <p>Range: 0 through 100</p> <p>Default: 100</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring PPPoE-Based Radio-to-Router Protocols on page 418

inline-jflow (Forwarding Options)

Syntax	<pre>inline-jflow { flow-export-rate <i>number</i>; source-address <i>ip-address</i>; }</pre>
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family inet output] [edit forwarding-options sampling instance <i>instance-name</i> family inet6 output]
Release Information	Statement introduced in Junos OS Release 10.4. Support for family inet6 added in Junos OS Release 12.1X45-D10.
Description	Specify Inline processing of sampled packets.
Options	<ul style="list-style-type: none">• flow-export-rate <i>value</i>—Flow export rate of monitored packets in kpps. The range is from 1 through 400.• source-address <i>address</i>—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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

interface (PIC Bundle)

Syntax	<pre>interface <i>interface-name</i>;</pre>
Hierarchy Level	[edit interfaces pic-set <i>pic-set-name</i>]
Release Information	Command introduced in Junos OS Release 9.6.
Description	Sets the PIC bundle and the interface.
Options	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

interface (PoE)

Syntax	<pre>interface (all <i>interface-name</i>) { disable; maximum-power <i>watts</i>; priority (high low); telemetries { disable; duration <i>hours</i>; interval <i>minutes</i>; } }</pre>
Hierarchy Level	[edit poe]
Release Information	Statement introduced in Junos OS Release 9.5.
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	<ul style="list-style-type: none"> • 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. <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Power over Ethernet on page 367

interfaces (CoS)

```
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*]

Release Information Statement introduced in Junos OS Release 8.5.

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.

Related Documentation • [Class of Service Feature Guide for Security Devices](#)

interval (Interfaces)

Syntax	<code>interval seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio-router credit]
Release Information	Statement introduced in Release 10.1 of Junos OS.
Description	Configure the frequency that the router generates credit announcement messages.
Options	<p>seconds—Interval between PADG credit announcements for each session.</p> <p>Range: 0 through 60</p> <p>Default: 1</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	• Configuring PPPoE-Based Radio-to-Router Protocols on page 418

interval (PoE)

Syntax	<code>interval minutes;</code>
Hierarchy Level	[edit poe interface (all <i>interface-name</i>) telemetries]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Modifies the interval for logging telemetries if you are monitoring the per-port power consumption for PoE interfaces.
Options	<p>minutes—Interval at which data is logged.</p> <p>Range: 1 through 30 minutes</p> <p>Default: 5 minutes</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	• Understanding Interfaces on page 3

ipv4-template (Services)

Syntax	ipv4-template;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Traffic Processing on Security Devices</i>• Understanding Interfaces on page 3

ipv6-template (Services)

Syntax	ipv6-template;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 12.1X45-D10.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Traffic Processing on Security Devices</i>• Understanding Interfaces on page 3

lcp (Interfaces)

Syntax	lcp { (active passive); periodic; }
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	For redundant Ethernet interfaces in a chassis cluster only, configure Link Aggregation Control Protocol (LACP).
Options	<ul style="list-style-type: none"> • active—Initiate transmission of LACP packets. • passive—Respond to LACP packets. • periodic—Interval for periodic transmission of LACP packets. <p>Default: If you do not specify lcp as either active or passive, LACP remains off (the default).</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding LACP on Standalone Devices on page 295 • periodic (Interfaces) on page 647

latency (Interfaces)

Syntax	latency <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio—router]
Release Information	Statement introduced in Junos OS Release 10.1.
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.
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 415

lease-time

Syntax	lease-time (<i>length</i> infinite);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a DHCP Client (CLI Procedure)</i>• <i>interfaces</i>• <i>unit</i>• <i>family</i>


line-rate (Interfaces)

Syntax	<code>line-rate</code>
Hierarchy Level	[edit interfaces <i>interfaces name</i> shdsl-options]
Release Information	Command introduced in Junos OS Release 10.0.
Description	Specify a line rate for an G.SHDSL interface.
Options	<ul style="list-style-type: none"> • 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.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the G.SHDSL Interface on SRX Series Devices</i>

link-speed (Interfaces)

Syntax	<code>link-speed <i>speed</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement modified in Release 9.0 of Junos OS.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Interfaces Configuration Guide for Security Devices</i>

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

Syntax	(loopback no-loopback);
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> gigether-options], [edit interfaces interface-range <i>name</i> ether-options]</p> <p>For QFX Series and EX Series:</p> <p>[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options],</p> <p>For SRX Series Devices and vSRX:</p> <p>[edit interfaces <i>interface-name</i> redundant-ether-options]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4 for MX Series.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Statement modified in Junos OS Release 9.2 for the SRX Series.</p>
Description	For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.
	<p> NOTE:</p> <ul style="list-style-type: none"> 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 <i>ge-fpc/pic/port</i> gigether-options] hierarchy level, an NDP address cannot be configured at the [edit interfaces <i>ge-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet6 address] hierarchy level.
Default	By default, loopback is disabled.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring Ethernet Loopback Capability](#)
 - [Understanding Interfaces on page 3](#)

loss-priority (CoS Loss Priority)

Syntax	<code>loss-priority <i>level</i> code-points [<i>values</i>];</code>
Hierarchy Level	[edit class-of-service loss-priority-maps frame-relay-de <i>map-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
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	<p><i>level</i> can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3 • Understanding Packet Loss Priorities

loss-priority (CoS Rewrite Rules)

Syntax	<code>loss-priority <i>level</i>;</code>
Hierarchy Level	[edit class-of-service rewrite-rules <i>type</i> <i>rewrite-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	<p>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.</p>
Options	<p><i>level</i> can be one of the following:</p> <ul style="list-style-type: none">• 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	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Class of Service Feature Guide for Security Devices</i>

loss-priority-maps (CoS Interfaces)

Syntax	<pre>loss-priority-maps { frame-relay-de (<i>map-name</i> default); }</pre>
Hierarchy Level	[edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Assign the loss priority map to a logical interface.
Options	<ul style="list-style-type: none"> • default—Apply default loss priority map. The default map contains the following: <pre>loss-priority low code-point 0; loss-priority high code-point 1;</pre> • 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.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

loss-priority-maps (CoS)

Syntax	<pre>loss-priority-maps { frame-relay-de <i>loss-priority-map-name</i> { loss-priority (high low medium-high medium-low) { code-points [<i>bit-string</i>]; } } }</pre>
Hierarchy Level	[edit class-of-service]
Release Information	Statement introduced in Junos OS Release 9.2.
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.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

management (PoE)

Syntax	management (class static);
Hierarchy Level	[edit poe]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Designates how the SRX Series device allocates power to the PoE ports.
Default	static
Options	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring PoE on All Interfaces on page 370

maximum-power (PoE)

Syntax	maximum-power <i>watts</i> ;
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 9.5.
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.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 370

media-type (Interfaces)

Syntax	media-type
Hierarchy Level	[edit interfaces <i>interface-name</i> media-type]
Release Information	Command introduced in Junos OS Release 10.2.
Description	Configure the operating modes for the 2-Port 10 Gigabit Ethernet XPIM.
Options	<ul style="list-style-type: none"> • copper • fiber
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

minimum-links (Interfaces)

Syntax	<code>minimum-links <i>number</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> redundant-ether-options]</code>
Release Information	Statement added in Release 10.1 of Junos OS.
Description	<p>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.</p> <p>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.</p>
Options	<p><i>number</i>—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.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Interfaces Configuration Guide for Security Devices</i>

mtu

Syntax	<code>mtu bytes;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
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	<i>bytes</i> —MTU size. Range: 0 through 5012 bytes Default: 1500 bytes (<i>inet</i> , <i>inet6</i> , and <i>iso</i> families), 1448 bytes (<i>mpls</i>)
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring MRRU on Multilink and Link Services Logical Interfaces Junos OS Network Interfaces Library for Routing Devices

native-vlan-id

Syntax	<code>native-vlan-id <i>vlan-id</i>;</code>
Hierarchy Level (QFX Series and EX4600)	<p>For platforms without ELS:</p> <p>[edit interfaces (QFX Series) <i>interface-name</i> unit 0 family ethernet-switching]</p> <p>For platforms with ELS:</p> <p>[edit interfaces (QFX Series) <i>interface-name</i>]</p>
Hierarchy Level (ACX Series, EX Series, SRX Series, M Series, MX Series, and T Series)	<p>[edit interfaces <i>ge-fpc/pic/port</i>],</p> <p>[edit interfaces <i>interface-name</i>]</p>
Hierarchy Level (SRX Series)	[edit interfaces <i>interface-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 9.5 for SRX Series.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>
Description	<p>Configure the VLAN identifier to associate with untagged packets received on the physical interface of a trunk mode interface for the following:</p> <ul style="list-style-type: none"> • 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 <p>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 <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p>

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.

Related Documentation	<ul style="list-style-type: none">• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>• <i>Configuring Gigabit Ethernet Interfaces (J-Web Procedure)</i>• <i>Understanding Bridging and VLANs on Switches</i>• <i>Enabling VLAN Tagging</i>• <i>Configuring Access Mode on a Logical Interface</i>• <i>Configuring the Native VLAN Identifier on Switches With ELS Support (CLI Procedure)</i>• Understanding Interfaces on page 3• <i>Understanding Q-in-Q Tunneling and VLAN Translation</i>• <i>no-native-vlan-insert</i>• <i>Sending Untagged Traffic Without VLAN ID to Remote End</i>• <i>show ethernet-switching interfaces</i>• <i>show vlans</i>• <i>flexible-vlan-tagging</i>• Junos OS Network Interfaces Configuration Guide
------------------------------	---

next-hop-tunnel

Syntax	<code>next-hop-tunnel gateway-address ipsec-vpn vpn-name;</code>
Hierarchy Level	<code>[edit interfaces interface-name unit logical-unit-number family family-name]</code>
Release Information	Statement introduced in Junos OS Release 9.5.
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	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

no-dns-propagation

Syntax	no-dns-propagation;
Hierarchy Level	[edit interface <i>interface-name</i> unit <i>unit-number</i> family <i>inet</i> <i>inet6 dhcp-client</i>]
Release Information	Statement introduced in Junos OS Release 12.1X47-D35.
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.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

option-refresh-rate (Services)

Syntax	option-refresh-rate
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the option refresh rate.
Options	<ul style="list-style-type: none"> • 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.
Related Documentation	<ul style="list-style-type: none"> • Configuring Flow Aggregation to Use Version 9 Flow Templates on page 32

pic-mode (Chassis T1 Mode)

Syntax	<code>pic-mode (clear-channel);</code>
Hierarchy Level	<code>[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> ethernet]</code>
Release Information	Statement added in Junos OS Release 10.2.
Description	Configure normal T1 mode or channelized T1 mode.
Options	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

periodic (Interfaces)

Syntax	<code>periodic (fast slow);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options lacp]
Release Information	Statement introduced in Junos OS Release 10.2.
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	<ul style="list-style-type: none">• fast—Transmit link aggregation control PDUs every second.• slow—Transmit link aggregation control PDUs every 30 seconds. <p>Default: <code>fast</code></p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 259

ppp-over-ether

Syntax	ppp-over-ether;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> encapsulation]
Release Information	Statement introduced before Junos OS Release 11.2. This encapsulation is supported for Redundant Ethernet interface in Junos OS Release 11.2.
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.
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 259

pppoe

Syntax	<pre>pppoe { command <i>binary-file-path</i>; disable; failover (alternate-media other-routing-engine); }</pre>
Hierarchy Level	[edit system processes]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Enable users to connect to a network of hosts over a bridge or access concentrator.
Options	<ul style="list-style-type: none"> • command <i>binary-file-path</i>—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. <ul style="list-style-type: none"> • 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 mastership 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.
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 259

pppoe-options

Syntax	<pre>pppoe-options { access-concentrator <i>name</i> ; auto-reconnect <i>seconds</i>; (client server); ignore-eol-tag; service-name <i>name</i>; underlying-interface <i>interface-name</i>; }</pre>
Hierarchy Level	[edit interfaces pp0 unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces pp0 unit <i>logical-unit-number</i>]
Release Information	Statement modified in Junos OS Release 12.3X48 to include ignore-eol-tag statement.
Description	Configure PPP over Ethernet-specific interface properties.
Options	<p>access-concentrator <i>name</i>—(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.</p> <p>auto-reconnect <i>seconds</i>—Configure the amount of time to wait before reconnecting after a session has terminated.</p> <p>client —Configure the device to operate in the PPPoE client mode.</p> <p>idle-timeout <i>seconds</i>—Configure the maximum time that a session can be idle.</p> <p>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.</p> <p>service-name <i>name</i>—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.</p> <p>server—Configure the device to operate in the PPPoE server mode.</p> <p>underlying-interface <i>interface-name</i>—Configure the interface on which PPP over Ethernet is running.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PPPoE Interfaces on page 393

priority (PoE)

Syntax	priority (high low);
Hierarchy Level	[edit poe interface (all <i>interface-name</i>)]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	<p>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.</p>
Default	low
Options	<p>value—high or low:</p> <ul style="list-style-type: none"> • 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	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 370

profile (Access)

```

Syntax  profile profile-name {
        accounting {
            accounting-stop-on-access-deny;
            accounting-stop-on-failure;
            coa-immediate-update;
            duplication;
            immediate-update;
            order [accounting-method];
            statistics (time | volume-time);
            update-interval minutes;
        }
        accounting-order [accounting-method];
        address-assignment pool pool-name;
        authentication-order [ldap | none | password | securid];
        authorization-order [jsrc];
        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;
        }
        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);

```

```

service {
  accounting-order {
    activation-protocol;
    radius;
  }
}
session-options {
  client-group [group-name];
  client-idle-timeout minutes;
  client-session-timeout minutes;
}
}

```

Hierarchy Level [edit access]

Release Information Statement introduced in Junos OS Release 10.4.

Description Create a profile containing a set of attributes that define device management access.

Required Privilege Level access—To view this statement in the configuration.
access-control—To add this statement to the configuration.

Related Documentation

- [Understanding Interfaces on page 3](#)
- *Understanding User Authentication for Security Devices*
- *Ethernet Switching and Layer 2 Transparent Mode Overview*

profiles

Syntax	<pre>profiles { profile-name { sip-user-id simple-ip-user-id; sip-password simple-ip-password; access-point-name apn; authentication-method (pap chap none); } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> cellular-options gsm-options]
Release Information	Statement introduced in Junos OS Release 9.5.
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	<p><i>profile-name</i>—Name of the profile.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 259

promiscuous-mode (Interfaces)

Syntax	<code>promiscuous-mode;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.1.
Description	<p>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.</p> <p>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.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling and Disabling Promiscuous Mode on Ethernet Interfaces (CLI Procedure) on page 270

quality (Interfaces)

Syntax	<code>quality <value>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> radio—router]</code>
Release Information	Statement introduced in Junos OS Release 10.1.
Description	This option controls relative link quality weight (value 0–100).
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 415

r2cp

Syntax	<pre>r2cp { command <i>binary-file-path</i>; disable; }</pre>
Hierarchy Level	[edit system processes]
Release Information	Statement introduced in Junos OS Release 9.2.
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	<ul style="list-style-type: none">• command <i>binary-file-path</i>—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.
Related Documentation	<ul style="list-style-type: none">• PPPoE-Based Radio-to-Router Protocols Overview on page 415

radio-router (Interfaces)

Syntax	<pre>radio-router { bandwidth <i>number</i>; credit { interval <i>number</i>; } data-rate <i>number</i>; latency <i>number</i>; quality <i>number</i>; resource <i>number</i>; threshold <i>number</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
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.
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 415

redundancy-group (Interfaces)

Syntax	<code>redundancy-group <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement introduced in Junos OS Release 9.0.
Description	Specify the redundancy group that a redundant Ethernet interface belongs to.
Options	<i>number</i> —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.
Related Documentation	<ul style="list-style-type: none">• Interfaces Feature Guide for Security Devices

redundant-ether-options

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*]

Release Information Statement introduced in Junos OS Release 9.2.

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.

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)

Syntax **redundant-parent** *interface-name*;

Hierarchy Level [edit interfaces *interface-name* fastether-options]

Release Information Statement introduced in Junos OS Release 9.2.

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.

Related Documentation

- [Understanding Ethernet Interfaces on page 259](#)

redundant-parent (Interfaces Gigabit Ethernet)

Syntax	<code>redundant-parent <i>interface-name</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> gigether-options]</code>
Release Information	Statement introduced in Release 9.0 of Junos OS.
Description	Configure Gigabit Ethernet-specific interface properties for Ethernet redundancy in a chassis cluster.
Options	<i>interface</i> —Parent redundant interface of the Gigabit Ethernet interface.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Interfaces Configuration Guide for Security Devices</i>

request pppoe connect

Syntax	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.
Description	Connect all sessions that are down.
Options	pppoe interface name — (Optional) Connect to a specified session.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• Understanding PPPoE Interfaces on page 392• Example: Configuring PPPoE Interfaces on page 393
List of Sample Output	request pppoe connect on page 662
Output Fields	When you enter this command, this command returns no output.

Sample Output

request pppoe connect

```
user@host> request pppoe connect
```

request pppoe disconnect

Syntax	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.
Description	Disconnect all active sessions.
Options	<p>session id — (Optional) Disconnect the session for which the session ID is specified.</p> <p>pppoe interface name— (Optional) Disconnect the session for a specific pppoe interface name.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • Understanding PPPoE Interfaces on page 392 • Example: Configuring PPPoE Interfaces on page 393
List of Sample Output	request pppoe disconnect on page 663
Output Fields	When you enter this command, this command returns no output.

Sample Output

request pppoe disconnect

```
user@host> request pppoe disconnect
```

resource (Interfaces)

Syntax	<code>resource <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> radio—router]
Release Information	Statement introduced in Junos OS Release 10.1.
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.
Related Documentation	<ul style="list-style-type: none">• PPPoE-Based Radio-to-Router Protocols Overview on page 415

retransmission-attempt

Syntax	<code>retransmission-attempt <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
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	<i>number</i> —Number of retransmit attempts. Range: 0 through 6 Default: 4
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a DHCP Client (CLI Procedure)• interfaces• unit• family

retransmission-interval (Interfaces)

Syntax	<code>retransmission-interval <i>seconds</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i> dhcp]</code>
Release Information	Statement introduced in Release 8.5 of Junos OS.
Description	Specify the time between successive retransmission attempts.
Options	<p><i>seconds</i> —Number of seconds between successive retransmission.</p> <p>Range: 4 through 64 seconds</p> <p>Default: 4 seconds</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Initial Configuration Guide for Security Devices</i>

roaming-mode

Syntax	<code>roaming-mode (home-only automatic)</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> cellular-options]</code>
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Specify whether the 3G wireless modem interface can access networks other than the home network.
Options	<ul style="list-style-type: none"> • home-only—No roaming is allowed. • automatic—Allows access to networks other than the home network. This is the default.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 259

scheduler-map (CoS Virtual Channels)

Syntax	<code>scheduler-map <i>map-name</i>;</code>
Hierarchy Level	<code>[edit class-of-service virtual-channel-groups <i>group-name</i> <i>virtual-channel-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Apply a scheduler map to this virtual channel.
Options	<i>map-name</i> —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.
Related Documentation	<ul style="list-style-type: none">• default (CoS)• shaping-rate (CoS Virtual Channels)• virtual-channel-group (CoS Interfaces)• virtual-channel-groups• virtual-channels

select-profile

Syntax	<code>select-profile <i>profile-name</i></code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> cellular-options gsm-options]</code>
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Select the active profile to establish a data call with a Global System for Mobile Communications (GSM) cellular network.
Options	<i>profile-name</i> —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.
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 259

server-address

Syntax	<code>server-address <i>ip-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
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	<i>ip-address</i> —DHCP server address.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a DHCP Client (CLI Procedure)</i> • <i>interfaces</i> • <i>unit</i> • <i>family</i>

shaping-rate (CoS Interfaces)

Syntax	<code>shaping-rate rate <overhead bytes> ;</code>
Hierarchy Level	<code>[edit class-of-service interfaces interface-name],</code> <code>[edit class-of-service interfaces interface-name unit logical-unit-number]</code>
Release Information	Statement introduced in Junos OS Release 9.2. overhead option introduced in Junos OS Release 18.1.
Description	<p>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.</p> <p>Logical and physical interface traffic shaping can be configured together. This means you can include the shaping-rate statement at the <code>[edit class-of-service interfaces interface interface-name]</code> hierarchy level <i>and</i> the <code>[edit class-of-service interfaces interface interface-name unit logical-unit-number]</code> 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.</p> <p>Alternatively, you can configure a shaping rate for a logical interface and oversubscribe the physical interface by including the shaping-rate statement at the <code>[edit class-of-service traffic-control-profiles]</code> hierarchy level. With this configuration approach, you can independently control the delay-buffer rate.</p> <p>On the physical interface, you can set the Layer 2 overhead adjustment to the shaping rate calculation at egress.</p>
Default	If you do not include this statement at the <code>[edit class-of-service interfaces interface interface-name unit logical-unit-number]</code> 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 <code>[edit class-of-service interfaces interface interface-name]</code> 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	<p>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</p> <p>overhead—Layer 2 shaping overhead adjustment to be applied at egress (bytes). Range: -62 through 192</p>
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.

Related Documentation • [policer-overhead](#)

simple-filter (Interfaces)

Syntax	<code>simple-filter;</code>
Hierarchy Level	<code>[edit interfaces <i>interfaces-name</i> unit <i>logical-unit-number</i> family <i>family-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Apply a simple filter to an interface. You can apply simple filters on ingress interfaces only.
Options	input <i>filter-name</i> : 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.
Related Documentation	• Understanding Ethernet Interfaces on page 259


sip-password

Syntax	<code>sip-password <i>simple-ip-password</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure the password provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.
Options	<i>simple-ip-password</i> —Password.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• Understanding Ethernet Interfaces on page 259

sip-user-id

Syntax	<code>sip-user-id <i>simple-ip-user-id</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> cellular-options gsm-options profiles <i>profile-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Configure the username provided by the service provider for connection to a Global System for Mobile Communications (GSM) cellular network.
Options	<i>simple-ip-user-id</i> —Username.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

source-address-filter (Interfaces)

Syntax	<code>source-address-filter <i>mac-address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement modified in Junos OS Release 9.2.
Description	<p>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.</p> <p>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.</p>
	<div>  <p>NOTE:</p> <ul style="list-style-type: none"> Software based MAC limiting is supported on SRX300, SRX320, and SRX340 devices. <p>A maximum of 32 devices are supported per device.</p> </div>
Options	<p><i>mac-address</i> —MAC address filter. You can specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nn:nn:nn:nn:nn:nn</i> (for example, 00:11:22:33:44:55) or <i>nnnn:nnnn:nnnn</i> (for example, 0011.2233.4455). You can configure up to 64 source addresses. To specify more than one address, include multiple <i>mac-address</i> options in the source-address-filter statement.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Understanding Ethernet Interfaces on page 259

source-filtering (Interfaces)

Syntax	(source-filtering no-source-filtering);
Hierarchy Level	[edit interfaces <i>interface-name</i> redundant-ether-options]
Release Information	Statement modified in Junos OS Release 9.2.
Description	<p>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.</p> <p>If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.</p> <p>By default, source address filtering is disabled.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 259

speed (Interfaces)

Syntax	speed (100m 10m 1g);
Hierarchy Level	[edit interfaces <i>interface-name</i> speed]
Release Information	Command introduced in Junos OS Release 10.2.
Description	Configure the operating speed for the 2-Port 10 Gigabit Ethernet XPIM.
Options	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• Understanding Ethernet Interfaces on page 259• Example: Configuring the 2-Port 10-Gigabit Ethernet XPIM Interface on page 328

speed (Gigabit Ethernet interface)

Syntax speed {1g |10g};

Hierarchy Level [edit interfaces *interface-name* gigether-options]

Release Information Statement introduced in Junos OS Release 18.1R1 for SRX4600.

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.

Following are the list of optics supported on SRX4600:

- SRX-SFP-IGE-LX
- SRX-SFP-IGE-LX-ET
- SRX-SFP-IGE-SX
- SRX-SFP-IGE-SX-ET
- SRX-SFP-IGE-T

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.
- 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 interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related • [speed \(Chassis Cluster\) on page 675](#)
Documentation

speed (Chassis Cluster)

Syntax speed {1g |10g};

Hierarchy Level [edit chassis cluster control-port]

Release Information Statement introduced in Junos OS Release 18.1R1 for SRX4600.

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.

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.

Related Documentation

- [speed \(Gigabit Ethernet interface\) on page 673](#)

telemetries (PoE)

Syntax

```
telemetries {  
    disable;  
    duration hours;  
    interval minutes;  
}
```

Hierarchy Level [edit poe interface (all | *interface-name*)]

Release Information Statement introduced in Junos OS Release 9.5.

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.

Related Documentation

- [Example: Configuring PoE on All Interfaces on page 370](#)

template-refresh-rate (Services)

Syntax	template-refresh-rate;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the template refresh rate.
Options	<ul style="list-style-type: none"> • 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.
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3

threshold (Interfaces)

Syntax	threshold <value>;
Hierarchy Level	[edit interfaces <i>interface-name</i> radio-router]
Release Information	Statement introduced in Junos OS Release 10.1.
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.
Related Documentation	<ul style="list-style-type: none"> • PPPoE-Based Radio-to-Router Protocols Overview on page 415

traceoptions (Interfaces)

Syntax	traceoptions
Hierarchy Level	[edit interfaces interface-name traceoptions]
Release Information	Command introduced in Junos OS Release 10.1.
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.
Related Documentation	<ul style="list-style-type: none">• PPPoE-Based Radio-to-Router Protocols Overview on page 415

update-server

Syntax	update-server;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet dhcp]
Release Information	Statement introduced in Junos OS Release 8.5 for J Series devices. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for SRX Series devices. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a DHCP Client (CLI Procedure)</i>• <i>Example: Configuring the Device as a DHCP Client</i>• <i>interfaces</i>• <i>unit</i>• <i>family</i>

vbr rate

Syntax	vbr rate;
Hierarchy Level	[edit interfaces interface-name atm-options vpi vpi-identifier shaping]
Release Information	Command introduced in Junos OS Release 9.5.
Description	For ATM encapsulation only, define a variable bit rate bandwidth utilization in the traffic-shaping profile.
Options	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

vdsl-profile

Syntax	vdsl-profile
Hierarchy Level	[edit interfaces interface-name vdsl-options]
Release Information	Command introduced in Junos OS Release 10.1.
Description	Configure the type of VDSL2 profiles. A profile is a table that contains a list of preconfigured VDSL2 settings.
Options	<ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• VDSL2 Interface Support on SRX Series Devices on page 184

vendor-id (Interfaces)

Syntax	<code>vendor-id <i>vendor-id</i> ;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i> dhcp]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Configure a vendor class ID for the Dynamic Host Configuration Protocol (DHCP) client.
Options	<i>vendor-id</i> —vendor class ID.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

web-authentication (Interfaces)

Syntax	<pre>web-authentication { http; https; redirect-to-https; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family-name</i> address <i>address</i>]
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, SRX550, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.
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.
Related Documentation	<ul style="list-style-type: none">• Understanding Interfaces on page 3

CHAPTER 39

Operational Commands

- clear oam ethernet connectivity-fault-management path-database
- clear dhcpv6 server binding (Local Server)
- clear ethernet-switching statistics mac-learning
- clear interfaces statistics swfabx
- clear ipv6 neighbors
- clear lacp statistics interfaces
- restart (Reset)
- request modem wireless create-profile
- request modem wireless fota
- request modem wireless sim-lock
- request modem wireless sim-unlock
- show chassis fpc (View)
- show chassis hardware (View)
- show ethernet-switching mac-learning-log
- show ethernet-switching table
- show igmp-snooping route (View)
- show interfaces
- show interfaces diagnostics optics
- show interfaces flow-statistics
- show interfaces queue
- show interfaces statistics (View)
- show interfaces terse zone
- show ipv6 neighbors
- show lacp interfaces (View)
- show lacp statistics interfaces (View)
- show modem wireless firmware
- show modem wireless network
- show modem wireless profiles

- [show oam ethernet link-fault-management](#)
- [show poe controller \(View\)](#)
- [show pppoe interfaces](#)
- [show pppoe statistics](#)
- [show poe telemetries](#)
- [show services accounting](#)
- [show services accounting aggregation \(View\)](#)
- [show services accounting aggregation template \(View\)](#)
- [show services accounting flow-detail \(View\)](#)

clear oam ethernet connectivity-fault-management path-database

Syntax	clear oam ethernet connectivity-fault-management path-database maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> host < <i>mac-addr</i> >
Release Information	Statement introduced in Junos OS Release 12.1X44-D10.
Description	Clear the relevant path information from the database for the specified remote host.
Options	<p>host—MAC address of remote host in xx:xx:xx:xx:xx:xx format.</p> <p>maintenance-association —Name of the maintenance association.</p> <p>maintenance-domain —Name of the maintenance domain.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • <i>show oam ethernet connectivity-fault-management path-database</i>
List of Sample Output	clear oam ethernet connectivity-fault- management path-database on page 685

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
```

clear dhcpv6 server binding (Local Server)

Syntax	<code>clear dhcpv6 server binding</code> <code><all <i>client-id</i> <i>ip-address</i> <i>session-id</i>></code> <code><interface <i>interface-name</i>></code> <code><routing-instance <i>routing-instance-name</i>></code>
Release Information	Command introduced in Junos OS Release 10.4.
Description	Clear the binding state of a DHCPv6 client from the client table on the DHCPv6 local server.
Options	<ul style="list-style-type: none">• <code>all</code>—(Optional) Clear the binding state for all DHCPv6 clients.• <code>client-id</code>—(Optional) Clear the binding state for the DHCPv6 client with the specified client ID (option 1).• <code>ip-address</code>—(Optional) Clear the binding state for the DHCPv6 client with the specified address.• <code>session-id</code>—(Optional) Clear the binding state for the DHCPv6 client with the specified session ID.• <code>interface <i>interface-name</i></code>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.• <code>routing-instance <i>routing-instance-name</i></code>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show dhcpv6 server binding (View)

clear ethernet-switching statistics mac-learning

Syntax	clear ethernet-switching statistics mac-learning
Release Information	Command introduced in Junos OS Release 10.1.
Description	Clear the media access control (MAC) learning statistics.
Options	<ul style="list-style-type: none">• none—Clear MAC learning statistics on all interfaces.• interface <i>interface-name</i>—(Optional) Clear MAC learning statistics on the specified interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ethernet-switching table on page 727show ethernet-switching table
List of Sample Output	clear ethernet-switching statistics mac-learning on page 687 clear ethernet-switching statistics mac-learning interface interface-name on page 687

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
```

clear interfaces statistics swfabx

Syntax	clear interfaces statistics <swfab0 swfab1>
Release Information	Command introduced in Junos OS Release 11.1.
Description	Clear interface statistics for the specified swfab interface.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• <i>show interfaces swfabx</i>
List of Sample Output	clear interfaces statistics <swfab0 swfab1> on page 688
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>
```


clear ipv6 neighbors

Syntax	clear ipv6 neighbors <all host <i>hostname</i> >
Release Information	Command introduced in Junos OS Release 12.1X45-D10.
Description	Clear IPv6 neighbor cache information.
Options	<p>none—Clear all IPv6 neighbor cache information.</p> <p>all—(Optional) Clear all IPv6 neighbor cache information.</p> <p>host <i>hostname</i>—(Optional) Clear the information for the specified IPv6 neighbors.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show ipv6 neighbors on page 845
List of Sample Output	clear ipv6 neighbors on page 689

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

```

clear lacp statistics interfaces

Syntax	<code>clear lacp statistics interfaces <interface-name></code>
Release Information	Command modified in Junos OS Release 10.2.
Description	Clear the LACP statistics. If you do not specify an interface name, LACP statistics for all interfaces are cleared.
Options	<i>interface-name</i> —(Optional) Name of an interface.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show lacp statistics interfaces (View) on page 851• Verifying LACP on Redundant Ethernet Interfaces on page 306
Output Fields	This command produces no output.

restart (Reset)

Syntax 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 | mntd-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 >

Release Information Command introduced before Junos OS Release 9.2

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 to drop calls and interrupt transmission, resulting in possible loss of data.

- Options**
- 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, for analyzing and tracking usage patterns, and for billing a user based upon the amount of time used or the type of services accessed.
 - 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.
 - commitd-service—(Optional) Restart the committed services.
 - database-replication—(Optional) Restart the database replication process.
 - datapath-trace-service—(Optional) Restart the Restart the packet path tracing process.

- `ddns`—(Optional) Restart the dynamic domain name system, which dynamically updates IP addresses for registered domain names.
- `dhcp`—(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.
- `disk-monitoring`—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.
- `dynamic-flow-capture`—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on PIC3 monitoring services cards.
- `ethernet-connectivity-fault-management`—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
- `ethernet-link-fault-management`—(Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
- `event-processing`—(Optional) Restart the event process (`eventd`).
- `extensible-subscriber-services`—(Optional) Restart the extensible subscriber services process.
- `fipsd`—(Optional) Restart the `fipsd` services.
- `firewall`—(Optional) Restart the firewall management process, which manages the firewall configuration and accepts or rejects packets that are transiting an interface on a router or switch.
- `firewall-authentication-service`—(Optional) Restart the firewall authentication service process.
- `general-authentication-service`—(Optional) Restart the general authentication process.
- `gprs-process`—(Optional) Restart the General Packet Radio Service (GPRS) process.
- `gracefully`—(Optional) Restart the software 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.
- `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-replication`—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.
- `lACP`—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link. The LACP process allows link aggregation control instances to reach agreement on the identity of the LAG to which a link belongs, moves the link to that LAG, and enables the transmission and reception processes for the link to function in an orderly manner.
- `l2cpd-service`—(SRX5400, SRX5600, and SRX5800 devices only) (Optional) Restart the Layer 2 Control Protocol (L2CP) process, which enables features such as L2 protocol tunneling and nonstop bridging.
- `l2-learning`—(Optional) Restart the Layer 2 (L2) address flooding and learning process.
- `license-service`—(Optional) Restart the feature license management process.
- `logical-system-service`—(Optional) Restart the logical system service process.
- `mib-process`—(Optional) Restart the MIB version II process, which provides the router's MIB II agent.
- `mountd-service`—(Optional) Restart the service for Network File System (NFS) mount requests.
- `named-service`—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.
- `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.
- `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 public key infrastructure (PKI) service process.
- `profillerd`—(Optional) Restart the profiler process.
- `remote-operations`—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.
- `rest-api`—(Optional) Restart the rest api process.
- `routing`—(Optional) Restart the routing protocol process (`rpd`).

- **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.
- **scc-chassisd**—(Optional) Restart the scc chassisd process.
- **secure-neighbor-discovery**—(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.
- **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.
- **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.
- **static-routed**—(Optional) Restart the static routed process.
- **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.
- **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 for L2 tunneled networks.
- **uac-service**—(Optional) Restart the Unified Access Control (UAC) process.
- **user-ad-authentication**—(Optional) Restart User ad Authentication process
- **vrrp**—(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**—(Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation • *Restart Commands Overview*

List of Sample Output [restart interfaces on page 695](#)

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
```

request modem wireless create-profile

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*

Release Information Command introduced in Junos OS 9.5. The **slot *sim-slot-number*** option is introduced in Junos OS 15.1X49-D100.

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

Related Documentation • [show modem wireless profiles on page 859](#)

List of Sample Output [request modem wireless create-profile on page 697](#)

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
```

request modem wireless fota

Syntax	<code>request modem wireless fota <i>interface-name</i> (enable disable)</code>
Release Information	Command introduced in Junos OS 15.1X49-D100.
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
Related Documentation	<ul style="list-style-type: none">• show modem wireless firmware on page 853
List of Sample Output	request modem wireless fota (enable) on page 698 request modem wireless fota (disable) on page 698

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
```

request modem wireless sim-lock

Syntax `request modem wireless sim-lock enable interface-name pin pin`

Release Information Command introduced in Junos OS Release 9.5.

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

Related Documentation

- [request modem wireless sim-unlock on page 700](#)

List of Sample Output [request modem wireless sim-lock on page 699](#)

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
```

request modem wireless sim-unlock

Syntax `request modem wireless sim-unlock interface-name pin unlock-code`

Release Information Command introduced in Junos OS Release 9.5.

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

Related Documentation

- [request modem wireless sim-lock on page 699](#)

List of Sample Output [request modem wireless sim-unlock on page 700](#)

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
```


show chassis fpc (View)

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)>>`

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 50 on page 702](#) summarizes the SRX5K-MPC3-40G10G (IOC3) PICs selected for various configuration scenarios.

Table 50: 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

Table 50: SRX5K-MPC3-40G10G (IOC3) PIC Selection Summary (continued)

CLI Configuration	PIC Selection
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
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.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

List of Sample Output [show chassis fpc on page 705](#)
[show chassis fpc \(SRX5600 and SRX5800 devices\) on page 705](#)
[show chassis fpc \(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G \(IOC3\) or SRX5K-MPC3-40G10G \(IOC3\) on page 705](#)
[show chassis fpc detail 2 on page 706](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices\) on page 706](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices with SPC2\) on page 706](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices with SRX5K-MPC\) on page 707](#)
[show chassis fpc pic-status \(SRX5600 and SRX5800 devices when Express Path \[formerly known as services offloading\] is configured\) on page 707](#)
[show chassis fpc pic-status \(with 20-Gigabit Ethernet MIC with SFP\) on page 708](#)
[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\) on page 708](#)
[show chassis fpc pic-status for HA \(SRX5600 and SRX5800 devices\) on page 708](#)
[show chassis fpc pic-status for HA\(SRX5400, SRX5600, and SRX5800 devices with SRX5K-MPC3-100G10G \(IOC3\) or SRX5K-MPC3-40G10G \(IOC3\) on page 709](#)

Output Fields [Table 51 on page 704](#) lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 51: 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: <ul style="list-style-type: none"> • 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.

Table 51: show chassis fpc Output Fields (continued)

Field Name	Field Description
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.
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 (%)
          (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 (%)
          (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
```

Slot	State	Temp (C)	CPU Utilization (%) Total	CPU Utilization (%) Interrupt	CPU Utilization (%) 1min	CPU Utilization (%) 5min	CPU Utilization (%) 15min	Memory DRAM (MB)
0	Online	36	20	Heap 0 Buffer 20	19	19	1024	
1	Online	35	8	4 0	26 8	8	8	2048
2	Online	40	21	12 0	14 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

```

show chassis hardware (View)

Syntax	show chassis hardware <clei-models detail extensive models node (<i>node-id</i> all local primary)>
Release Information	Command introduced in Junos OS Release 9.2. Command modified in Junos OS Release 9.2 to include node option.
Description	Display chassis hardware information.
Options	<ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • 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
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Traffic Processing on Security Devices</i> • Interface Naming Conventions on page 8
Output Fields	Table 52 on page 710 lists the output fields for the show chassis hardware command. Output fields are listed in the approximate order in which they appear.

Table 52: 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.
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.

Table 52: show chassis hardware Output Fields (continued)

Field Name	Field Description
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 52: show chassis hardware Output Fields (continued)

Field Name	Field Description
Description	<p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> Type of power supply. Switch Control Board (SCB) <p>Starting with Junos OS Release 12.1X47-D15 and Junos OS Release 17.3R1, the SRX5K-SCBE (SCB2) is introduced.</p> <ul style="list-style-type: none"> 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. <p>Starting with Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1, the SRX5K-SCB3 (SCB3) with enhanced midplane is introduced.</p> <ul style="list-style-type: none"> All existing SCB software that is supported by SCB2 is supported on SCB3. SRX5K-RE-1800X4 (RE2). Mixed Routing Engine use is not supported. 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 master 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.

Table 52: show chassis hardware Output Fields (continued)

Field Name	Field Description
	<ul style="list-style-type: none"> Type of Flexible PIC Concentrator (FPC), Physical Interface Card (PIC), Modular Interface Cards (MICs), and PIMs. IOCs <p>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.</p> <ul style="list-style-type: none"> 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. IOC3 supports intra-chassis and inter-chassis fabric redundancy mode. IOC3 supports ISSU and ISHU in chassis cluster mode. IOC3 supports intra-FPC and Inter-FPC Express Path (previously known as <i>services offloading</i>) 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. <p>NOTE: Fabric bandwidth increasing mode is not supported on IOC3.</p> SRX Clustering Module (SCM) Fan tray For hosts, the Routing Engine type. <ul style="list-style-type: none"> Starting with Junos OS Release 12.1X47-D15 and Junos OS Release 17.3R1, the SRX5K-RE-1800X4 (RE2) Routing Engine is introduced. The RE2 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. <p>NOTE: The RE2 provides significantly better performance than the previously used Routing Engine, even with a single core.</p>

show chassis hardware

show chassis hardware

```
user@host> show chassis hardware
node0:
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B190FAGB	SRX5600
Midplane	REV 01	710-024804	ABAB5282	SRX5600 Midplane
FPM Board	REV 01	710-024631	YG0211	Front Panel Display
PEM 0	Rev 03	740-034701	QCS13090901H	PS 1.4-2.6kW; 90-264V
AC in				
PEM 1	Rev 02	740-034701	QCS130309005	PS 1.4-2.6kW; 90-264V
AC in				
PEM 2	Rev 01	740-034701	QCS12190901A	PS 1.4-2.6kW; 90-264V
AC in				
Routing Engine 0	REV 01	740-056658	9009150226	SRX5k RE-1800X4
CB 0	REV 05	750-066337	CAJS6543	SRX5k SCB3
Xcvr 0				
FPC 0	REV 01	750-077373	CAKC4112	SPC3
CPU		BUILTIN	BUILTIN	SRX5k vCPP Broadwell
FPC 5	REV 08	750-043157	CABL8327	SRX5k IOC II
CPU	REV 03	711-043360	CABJ0770	SRX5k MPC PMB
MIC 1	REV 01	750-055732	ZM8169	20x 1GE(LAN) SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 02	740-013111	A514696	SFP-T
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-013111	70173027	SFP-T
Xcvr 9	REV 01	740-030658	AD1130A01S2	UNSUPPORTED
Fan Tray				Enhanced Fan Tray

node1:

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11ADC7CAGB	SRX5600
Midplane	REV 01	710-024804	ABAB5357	SRX5600 Midplane
FPM Board	REV 01	710-024631	YF2474	Front Panel Display
PEM 0	Rev 03	740-034701	QCS133309019	PS 1.4-2.6kW; 90-264V
AC in				
PEM 1	Rev 03	740-034701	QCS133209023	PS 1.4-2.6kW; 90-264V
AC in				
PEM 2	Rev 02	740-034701	QCS130309013	PS 1.4-2.6kW; 90-264V
AC in				
Routing Engine 0	REV 03	740-049603	9013079642	RE-S-EX9200-1800X4
CB 0	REV 01	750-056587	CACC9541	SRX5k SCB II
Xcvr 0				
FPC 0	REV 01	750-077373	CAKC4097	SPC3
CPU		BUILTIN	BUILTIN	SRX5k vCPP Broadwell
FPC 5	REV 11	750-043157	CACA8792	SRX5k IOC II
CPU	REV 04	711-043360	CACA8809	SRX5k MPC PMB
MIC 1	REV 01	750-055732	CACF9067	20x 1GE(LAN) SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-013111	8512082	SFP-T
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-013111	70113020	SFP-T
Xcvr 9	REV 01	740-021308	AJQ058Y	UNSUPPORTED
Fan Tray				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 Module
PEM 0	Rev 03	740-034724	QCS17020203F	PS 4.1kW; 200-240V AC i
PEM 1	Rev 03	740-034724	QCS17020203C	PS 4.1kW; 200-240V AC i
PEM 2	Rev 04	740-034724	QCS17100200A	PS 4.1kW; 200-240V AC i
PEM 3	Rev 03	740-034724	QCS17080200M	PS 4.1kW; 200-240V AC i
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	AMBOHG3	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	MUCOMF2	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
PEM 1	Rev 03	740-034701	QCS13090904T	PS 1.4-2.6kW; 90-264V A
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			JN124FECOAGB	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)

```
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	ALNON6C	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 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

show ethernet-switching mac-learning-log

Syntax	show ethernet-switching mac-learning-log
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 devices. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Displays the event log of learned MAC addresses.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ethernet-switching table on page 727 • <i>show ethernet-switching interfaces</i> • show ethernet-switching table on page 727 • <i>show ethernet-switching interfaces</i> • <i>Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch</i> • <i>Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</i> • <i>Example: Connecting an EX Series Access Switch to a Distribution Switch</i>
List of Sample Output	show ethernet-switching mac-learning-log (EX Series switch) on page 724 show ethernet-switching mac-learning-log (QFX Series Switches, QFabric, NFX Series Devices and EX4600) on page 724 show ethernet-switching mac-learning-log (SRX Series devices) on page 725
Output Fields	<p>Output fields for EX Series switches:</p> <p>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.</p>

Table 53: 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.

Table 53: show ethernet-switching mac-learning-log Output Fields (continued)

Field Name	Field Description
Blocking	The forwarding state of the interface: <ul style="list-style-type: none"> • 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 54 on page 723](#) 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 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.
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 55 on page 723](#) 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 55: 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.

Table 55: show ethernet-switching-mac-learning-log Output Fields (continued)

Field Name	Field Description
Deleted Added	MAC address deleted or added to the MAC learning log.
Blocking	<p>The forwarding state of the interface:</p> <ul style="list-style-type: none"> 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]
```

show ethernet-switching table

List of Syntax	Syntax (QFX Series, QFabric, NFX Series and EX4600) on page 727 Syntax (EX Series) on page 727 Syntax (EX Series, MX Series and QFX Series) on page 727 Syntax (SRX Series) on page 727
Syntax (QFX Series, QFabric, NFX Series and EX4600)	<pre>show ethernet-switching table <brief detail extensive summary> <interface <i>interface-name</i>> <management-vlan> <sort-by (<i>name</i> <i>tag</i>)> <vlan <i>vlan-name</i>></pre>
Syntax (EX Series)	<pre>show ethernet-switching table <brief detail extensive summary> <interface <i>interface-name</i>> <management-vlan> <persistent-mac <interface <i>interface-name</i>>> <sort-by (<i>name</i> <i>tag</i>)> <vlan <i>vlan-name</i>></pre>
Syntax (EX Series, MX Series and QFX Series)	<pre>show ethernet-switching table <brief count detail extensive summary> <address> <instance <i>instance-name</i>> <interface <i>interface-name</i>> isis <i>isid</i> <logical-system <i>logical-system-name</i>> <persistent-learning (interface <i>interface-name</i> mac <i>mac-address</i>)> <address> <vlan-id (all-vlan <i>vlan-id</i>)> <vlan-name (all <i>vlan-name</i>)></pre>
Syntax (SRX Series)	<pre>show ethernet-switching table (brief detail extensive) interface <i>interface-name</i></pre>
Release Information	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 9.5 for SRX Series.</p> <p>Options summary, management-vlan, and vlan <i>vlan-name</i> introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Option sort-by and field name tag introduced in Junos OS Release 10.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Output for private VLANs introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Option persistent-mac introduced in Junos OS Release 11.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Options logical-system, persistent-learning, and summary introduced in Junos OS Release 13.2X50-D10 (ELS).</p>

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

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*
- *clear ethernet-switching table*
- [show ethernet-switching mac-learning-log on page 722](#)

List of Sample Output [show ethernet-switching table \(Enhanced Layer 2 Software on QFX Series, QFabric, NFX Series and EX460\) on page 733](#)
[show ethernet-switching table \(QFX Series, QFabric, NFX Series and EX460\) on page 734](#)
[show ethernet-switching table \(Private VLANs on QFX Series, QFabric, NFX Series and EX460\) on page 735](#)
[show ethernet-switching table \(Junos Fusion Data Center with EVPN on QFX Series switches\) on page 735](#)
[show ethernet-switching table brief \(QFX Series, QFabric, NFX Series and EX460\) on page 736](#)
[show ethernet-switching table detail \(QFX Series, QFabric, NFX Series and EX460\) on page 737](#)
[show ethernet-switching table extensive \(QFX Series, QFabric, NFX Series and EX460\) on page 738](#)
[show ethernet-switching table interface \(QFX Series, QFabric, NFX Series and EX460\) on page 740](#)
[show ethernet-switching table \(EX Series switches\) on page 740](#)
[show ethernet-switching table brief \(EX Series switches\) on page 740](#)
[show ethernet-switching table detail \(EX Series switches\) on page 741](#)

[show ethernet-switching table extensive \(EX Series switches\) on page 742](#)
[show ethernet-switching table persistent-mac \(EX Series switches\) on page 742](#)
[show ethernet-switching table persistent-mac interface ge-0/0/16.0 \(EX Series switches\) on page 742](#)
[show ethernet-switching table \(EX Series, MX Series and QFX Series\) on page 742](#)
[show ethernet-switching table brief on page 744](#)
[show ethernet-switching table count on page 745](#)
[show ethernet-switching table extensive on page 746](#)
[show ethernet-switching table detail \(SRX Series\) on page 747](#)
[show ethernet-switching table extensive \(SRX Series\) on page 748](#)
[show ethernet-switching table interface ge-0/0/1 \(SRX Series\) on page 749](#)

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 56: show ethernet-switching table Output Fields

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: <ul style="list-style-type: none"> • 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
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 57: *show ethernet-switching table* Output Fields

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: <ul style="list-style-type: none"> • 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
Type	The type of MAC address. Values are: <ul style="list-style-type: none"> • 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 58: *show ethernet-switching table* Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.
VLAN name	Name of the VLAN.

Table 58: show ethernet-switching table Output fields (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: <ul style="list-style-type: none"> • 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.
Age	This field is not supported.
Logical interface	Name of the logical interface.
Active source	IP address of remote entity on which MAC address is learned.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
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-protocol epoch 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 59 on page 732](#) lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 59: show ethernet-switching table Output Fields

Field Name	Field Description
VLAN	The name of a VLAN.

Table 59: show ethernet-switching table Output Fields (continued)

Field Name	Field Description
MAC address	The MAC address associated with the VLAN.
Type	The type of MAC address. Values are: <ul style="list-style-type: none"> 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,
0 - ovsdb MAC)
```

```
Ethernet switching table : 2 entries, 2 learned
```

```
Routing instance : default-switch
```

Vlan name	MAC address	MAC flags	Age	Logical 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,
0 - ovsdb MAC)
```

```
Ethernet switching table : 2 entries, 2 learned
```

```
Routing instance : default-switch
```

Vlan name	MAC address	MAC flags	Age	Logical 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,
```

0 - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical 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,
0 - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical 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,
0 - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical 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, QFabric, NFX Series and EX460)

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 (Junos Fusion Data Center with EVPN on QFX Series switches)

```

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 : 30 entries, 30 learned
Routing instance : default-switch
Vlan      MAC      MAC      Logical      Active
name      address   flags    interface    source
v100      00:31:46:e8:f9:d6 D        vtep.32768
192.168.2.22
v100      7c:e2:ca:e2:75:7c D        vtep.32771
192.168.4.44
v100      7c:e2:ca:e4:05:9a D        vtep.32770
192.168.3.33
v101      00:31:46:e8:f9:d6 D        vtep.32768
192.168.2.22
v101      7c:e2:ca:e2:75:7c D        vtep.32771
192.168.4.44
v101      7c:e2:ca:e4:05:9a D        vtep.32770
192.168.3.33
v102      00:31:46:e8:f9:d6 D        vtep.32768

```

192.168.2.22				
v102	7c:e2:ca:e2:75:7c	D		vtep.32771
192.168.4.44				
v102	7c:e2:ca:e4:05:9a	D		vtep.32770
192.168.3.33				
v103	00:31:46:e8:f9:d6	D		vtep.32768
192.168.2.22				
v103	7c:e2:ca:e2:75:7c	D		vtep.32771
192.168.4.44				
v103	7c:e2:ca:e4:05:9a	D		vtep.32770
192.168.3.33				
v3001	00:31:46:e8:f9:d6	D		vtep.32768
192.168.2.22				
v3001	28:c0:da:6a:9f:c2	DL		ae11.0
v3001	7c:e2:ca:e2:75:7c	D		vtep.32771
192.168.4.44				
v3001	7c:e2:ca:e4:05:9a	D		vtep.32770
192.168.3.33				
v3002	00:31:46:e8:f9:d6	D		vtep.32768
192.168.2.22				
v3002	7c:e2:ca:e2:75:7c	D		vtep.32771
192.168.4.44				
v3002	7c:e2:ca:e4:05:9a	D		vtep.32770
192.168.3.33				
v3003	00:31:46:e8:f9:d6	D		vtep.32768
192.168.2.22				
v3003	28:c0:da:6a:9f:c2	DL		ae11.0
v3003	7c:e2:ca:e2:75:7c	D		vtep.32771
192.168.4.44				
v3003	7c:e2:ca:e4:05:9a	D		vtep.32770
192.168.3.33				
v3004	00:31:46:e8:f9:d6	D		vtep.32768
192.168.2.22				
v3004	7c:e2:ca:e2:75:7c	D		vtep.32771
192.168.4.44				
v3004	7c:e2:ca:e4:05:9a	D		vtep.32770
192.168.3.33				
v3005	00:31:46:e8:f9:d6	D		vtep.32768
192.168.2.22				
v3005	28:c0:da:6a:9f:c2	DL		ae11.0
v3005	7c:e2:ca:e2:75:7c	D		vtep.32771
192.168.4.44				
v3005	7c:e2:ca:e4:05:9a	D		vtep.32770
192.168.3.33				

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
```

show igmp-snooping route (View)

Syntax	show igmp-snooping route (brief detail ethernet-switching inet vlan)
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display IGMP snooping route information.
Options	<ul style="list-style-type: none"> • 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 <i>vlan-id</i> <i>vlan-name</i>—(Optional) Display route information for the specified VLAN.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3
Output Fields	Table 60 on page 750 lists the output fields for the show igmp-snooping route command. Output fields are listed in the approximate order in which they appear.

Table 60: 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		

show interfaces

List of Syntax [Syntax \(Gigabit Ethernet\) on page 752](#)
 [Syntax \(10 Gigabit Ethernet\) on page 752](#)
 [Syntax \(SRX Series Devices\) on page 752](#)

Syntax (Gigabit Ethernet) `show interfaces ge-fpc/pic/port`
 `<brief | detail | extensive | terse>`
 `<descriptions>`
 `<media>`
 `<snmp-index snmp-index>`
 `<statistics>`

Syntax (10 Gigabit Ethernet) `show interfaces xe-fpc/pic/port`
 `<brief | detail | extensive | terse>`
 `<descriptions>`
 `<media>`
 `<snmp-index snmp-index>`
 `<statistics>`

Syntax (SRX Series Devices) `show interfaces (`
 `<interface-name>`
 `<brief | detail | extensive | terse>`
 `<controller interface-name>|`
 `<descriptions interface-name>|`
 `<destination-class (all | destination-class-name logical-interface-name)>|`
 `<diagnostics optics interface-name>|`
 `<far-end-interval interface-fpc/pic/port>|`
 `<filters interface-name>|`
 `<flow-statistics interface-name>|`
 `<interval interface-name>|`
 `<load-balancing (detail | interface-name)>|`
 `<mac-database mac-address mac-address>|`
 `<mc-ae id identifier unit number revertive-info>|`
 `<media interface-name>|`
 `<policers interface-name>|`
 `<queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics>|`
 `<redundancy (detail | interface-name)>|`
 `<routing brief detail summary interface-name>|`
 `<routing-instance (all | instance-name)>|`
 `<snmp-index snmp-index>|`
 `<source-class (all | destination-class-name logical-interface-name)>|`
 `<statistics interface-name>|`
 `<switch-port switch-port number>|`
 `<transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all |`
 `interface-name)>|`
 `<zone interface-name>`
 `)`

Release Information Command introduced before Junos OS Release 7.4 for Gigabit interfaces.
 Command introduced in Junos OS Release 8.0 for 10 Gigabit interfaces.
 Command modified in Junos OS Release 9.5 for SRX Series devices.

Command introduced in Junos OS Release 18.1 for Gigabit interfaces.

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 appliance running Junos OS.



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

Options For Gigabit interfaces:

ge-fpc/pic/port—Display standard information about the specified 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 10 Gigabit interfaces:

xe-fpc/pic/port—Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

For SRX interfaces:

- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and port with the port number.
 - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
 - **ce1-*pim*/0/*port***—Channelized E1 interface.
 - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
 - **ct1-*pim*/0/*port***—Channelized T1 interface.
 - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
 - **e1-*pim*/0/*port***—E1 interface.
 - **e3-*pim*/0/*port***—E3 interface.
 - **fe-*pim*/0/*port***—Fast Ethernet interface.
 - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
 - **se-*pim*/0/*port***—Serial interface.
 - **t1-*pim*/0/*port***—T1 (also called DS1) interface.
 - **t3-*pim*/0/*port***—T3 (also called DS3) interface.
 - **wx-slot/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace *pim* with the PIM slot and port with the port number.
 - **at-*pim*/0/*port***—ATM-over-ADSL or ATM-over-SHDSL interface.
 - **ce1-*pim*/0/*port***—Channelized E1 interface.
 - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
 - **ct1-*pim*/0/*port***—Channelized T1 interface.
 - **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
 - **e1-*pim*/0/*port***—E1 interface.
 - **e3-*pim*/0/*port***—E3 interface.
 - **fe-*pim*/0/*port***—Fast Ethernet interface.
 - **ge-*pim*/0/*port***—Gigabit Ethernet interface.
 - **se-*pim*/0/*port***—Serial interface.
 - **t1-*pim*/0/*port***—T1 (also called DS1) interface.

- **t3-pim/0/port**—T3 (also called DS3) interface.
- **wx-slot/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).

Additional Information In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.

Required Privilege Level view

Related Documentation

- *Understanding Layer 2 Interfaces on Security Devices*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- *Verifying and Managing Configurations for Dynamic VLANs Based on Access-Line Identifiers*

List of Sample Output

[show interfaces \(Gigabit Ethernet\) on page 792](#)
[show interfaces \(Gigabit Ethernet on MX Series Routers\) on page 792](#)
[show interfaces \(link degrade status\) on page 793](#)
[show interfaces extensive \(Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration\) on page 793](#)
[show interfaces brief \(Gigabit Ethernet\) on page 794](#)
[show interfaces detail \(Gigabit Ethernet\) on page 794](#)
[show interfaces extensive \(Gigabit Ethernet IQ2\) on page 796](#)
[show interfaces \(Gigabit Ethernet Unnumbered Interface\) on page 799](#)
[show interfaces \(ACI Interface Set Configured\) on page 799](#)
[show interfaces \(ALI Interface Set\) on page 799](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, IQ2\) on page 800](#)
[show interfaces extensive \(10-Gigabit Ethernet, WAN PHY Mode\) on page 802](#)
[show interfaces extensive \(10-Gigabit Ethernet, DWDM OTN PIC\) on page 804](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode\) on page 806](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only\) on page 807](#)
[show interfaces extensive \(10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only\) on page 808](#)
[Sample Output SRX Gigabit Ethernet on page 809](#)
[Sample Output SRX Gigabit Ethernet on page 809](#)
[show interfaces detail \(Gigabit Ethernet\) on page 810](#)
[show interfaces statistics st0.0 detail on page 812](#)
[show interfaces extensive \(Gigabit Ethernet\) on page 813](#)
[show interfaces terse on page 815](#)
[show interfaces controller \(Channelized E1 IQ with Logical E1\) on page 816](#)
[show interfaces controller \(Channelized E1 IQ with Logical DSO\) on page 816](#)
[show interfaces descriptions on page 816](#)
[show interfaces destination-class all on page 816](#)

[show interfaces diagnostics optics](#) on page 817
[show interfaces far-end-interval coc12-5/2/0](#) on page 817
[show interfaces far-end-interval coc1-5/2/1:1](#) on page 818
[show interfaces filters](#) on page 818
[show interfaces flow-statistics \(Gigabit Ethernet\)](#) on page 818
[show interfaces interval \(Channelized OC12\)](#) on page 819
[show interfaces interval \(E3\)](#) on page 820
[show interfaces interval \(SONET/SDH\) \(SRX devices\)](#) on page 820
[show interfaces load-balancing \(SRX devices\)](#) on page 820
[show interfaces load-balancing detail \(SRX devices\)](#) on page 821
[show interfaces mac-database \(All MAC Addresses on a Port SRX devices\)](#) on page 821
[show interfaces mac-database \(All MAC Addresses on a Service SRX devices\)](#) on page 821
[show interfaces mac-database mac-address](#) on page 822
[show interfaces mc-ae \(SRX devices\)](#) on page 822
[show interfaces media \(SONET/SDH\)](#) on page 822
[show interfaces policers \(SRX devices\)](#) on page 823
[show interfaces policers interface-name \(SRX devices\)](#) on page 823
[show interfaces queue \(SRX devices\)](#) on page 823
[show interfaces redundancy \(SRX devices\)](#) on page 824
[show interfaces redundancy \(Aggregated Ethernet SRX devices\)](#) on page 825
[show interfaces redundancy detail \(SRX devices\)](#) on page 825
[show interfaces routing brief \(SRX devices\)](#) on page 825
[show interfaces routing detail \(SRX devices\)](#) on page 825
[show interfaces routing-instance all \(SRX devices\)](#) on page 826
[show interfaces snmp-index \(SRX devices\)](#) on page 826
[show interfaces source-class all \(SRX devices\)](#) on page 826
[show interfaces statistics \(Fast Ethernet SRX devices\)](#) on page 827
[show interfaces switch-port \(SRX devices\)](#) on page 827
[show interfaces transport pm \(SRX devices\)](#) on page 828
[show security zones \(SRX devices\)](#) on page 829

Output Fields [Table 61 on page 756](#) 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 62 on page 784](#).

Table 61: 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 <i>Common Output Fields Description</i> .	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

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under <i>Common Output Fields Description</i> .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: 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
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	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces command.</p>	detail extensive

Table 61: *show interfaces (Gigabit Ethernet) Output Fields (continued)*

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Drops field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number must always be 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	<p>Total number of egress queues supported on the specified interface.</p> <p>NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GER, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the show interfaces command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenario, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs</p>	detail extensive

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Dropped packets field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Interface transmit statistics	<p>(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled.</p> <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—Count of corrected errors in the last second. • Corrected Error Ratio—Corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—Number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errored blocks—Number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive
Link Degrad	<p>Shows the link degrade status of the physical link and the estimated bit error rates (BERs). This field is available only for the PICs supporting the physical link monitoring feature.</p> <ul style="list-style-type: none"> • Link Monitoring—Indicates if physical link degrade monitoring is enabled on the interface. <ul style="list-style-type: none"> • Enable—Indicates that link degrade monitoring has been enabled (using the <code>link-degrade-monitor</code> 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 Degrad Set Threshold—The BER threshold value at which the link is considered degraded and a corrective action is triggered. • Link Degrad Clear Threshold—The BER threshold value at which the degraded link is considered recovered and the corrective action applied to the interface is reverted. • Estimated BER—The estimated bit error rate. • Link-degrade event—Shows link degrade event information. <ul style="list-style-type: none"> • Seconds—Time (in seconds) elapsed after a link degrade event occurred. • Count—The number of link degrade events recorded. • State—Shows the link degrade status (example: Defect Active). 	detail extensive

Table 61: *show interfaces (Gigabit Ethernet) Output Fields (continued)*

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces command. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. <p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> • 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

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet may enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field must increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field must not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner—Information from the remote Ethernet device: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the link partner, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), Symmetric/Asymmetric (link partner supports PAUSE on receive and transmit or only PAUSE on transmit), and None (link partner does not support flow control). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the local Ethernet device: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are Symmetric/Asymmetric (local device supports PAUSE on receive and transmit or only PAUSE on receive) and None (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display Symmetric (local device supports PAUSE on receive and transmit), Asymmetric (local device supports PAUSE on receive), and None (local device does not support flow control). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	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 <i>Common Output Fields Description</i> .	All levels
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Demux	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN	<p>Information displayed for agent circuit identifier (ACI) interface set configured with the agent-circuit-id autoconfiguration stanza.</p> <p>Dynamic Profile—Name of the dynamic profile that defines the ACI interface set.</p> <p>If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.</p> <p>NOTE: The ACI VLAN field is replaced with the Line Identity field when an ALI interface set is configured with the line-identity autoconfiguration stanza.</p>	brief detail extensive none
Line Identity	<p>Information displayed for access-line-identifier (ALI) interface sets configured with the line-identity autoconfiguration stanza.</p> <ul style="list-style-type: none"> • Dynamic Profile—Name of the dynamic profile that defines the ALI interface set. • Trusted option used to create the ALI interface set: Circuit-id, Remote-id, or Accept-no-ids. More than one option can be configured. <p>If configured, the ALI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ALI information.</p> <p>NOTE: The Line Identity field is replaced with the ACI VLAN field when an ACI interface set is configured with the agent-circuit-id autoconfiguration stanza.</p>	detail
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Neighbor Discovery Protocol (NDP) Queue Statistics	<p>NDP statistics for protocol inet6 under logical interface statistics.</p> <ul style="list-style-type: none"> • Max nh cache—Maximum interface neighbor discovery nexthop cache size. • New hold nh limit—Maximum number of new unresolved nexthops. • Curr nh cnt—Current number of resolved nexthops in the NDP queue. • Curr new hold cnt—Current number of unresolved nexthops in the NDP queue. • NH drop cnt—Number of NDP requests not serviced. 	All levels
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

Table 61: show interfaces (Gigabit Ethernet) Output Fields (continued)

Field Name	Field Description	Level of Output
Max Sessions	Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none
Duplicate Protection	State of PPPoE duplicate protection: On or Off . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	detail extensive none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: On or Off . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	detail extensive none
AC Name	Name of the access concentrator.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive

Table 61: *show interfaces (Gigabit Ethernet) Output Fields (continued)*

Field Name	Field Description	Level of Output
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 <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag. Possible values are described in the "Addresses Flags" section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

The following table describes the output fields for the **show interfaces** (10-Gigabit Ethernet) command.

Field Name	Field Description	Level of Output
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none

SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link flags	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none

Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive

Egress queues	Total number of egress queues supported on the specified interface. 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. <ul style="list-style-type: none">• Queued packets—Number of queued packets.• Transmitted packets—Number of transmitted packets.• Dropped packets—Number of packets dropped by the ASIC's RED mechanism.	detail extensive
Ingress queues	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. <ul style="list-style-type: none">• Queued packets—Number of queued packets.• Transmitted packets—Number of transmitted packets.• Dropped packets—Number of packets dropped by the ASIC's RED mechanism.	extensive
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the 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 . <ul style="list-style-type: none">• None—There are no active defects or alarms.• Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	detail extensive none
OTN alarms	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. <ul style="list-style-type: none">• efec—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors.• gfec—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors.• none—FEC mode is not configured.	detail extensive

OTN Rate	<p>OTN mode.</p> <ul style="list-style-type: none"> • fixed-stuff-bytes—Fixed stuff bytes 11.0957 Gbps. • no-fixed-stuff-bytes—No fixed stuff bytes 11.0491 Gbps. • pass-through—Enable OTN passthrough mode. • no-pass-through—Do not enable OTN passthrough mode. 	detail extensive
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: enabled or disabled .	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
OTN FEC alarms	<p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <ul style="list-style-type: none"> • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect. 	detail extensive
OTN OC	<p>OTN OC defects triggered on the interface.</p> <ul style="list-style-type: none"> • LOS—OC Loss of Signal defect. • LOF—OC Loss of Frame defect. • LOM—OC Loss of Multiframe defect. • Wavelength Lock—OC Wavelength Lock defect. 	detail extensive
OTN OTU	<p>OTN OTU defects detected on the interface</p> <ul style="list-style-type: none"> • AIS—OTN AIS alarm. • BDI—OTN OTU BDI alarm. • IAE—OTN OTU IAE alarm. • TTIM—OTN OTU TTIM alarm. • SF—OTN ODU bit error rate fault alarm. • SD—OTN ODU bit error rate defect alarm. • TCA-ES—OTN ODU ES threshold alarm. • TCA-SES—OTN ODU SES threshold alarm. • TCA-UAS—OTN ODU UAS threshold alarm. • TCA-BBE—OTN ODU BBE threshold alarm. • BIP—OTN ODU BIP threshold alarm. • BBE—OTN OTU BBE threshold alarm. • ES—OTN OTU ES threshold alarm. • SES—OTN OTU SES threshold alarm. • UAS—OTN OTU UAS threshold alarm. 	detail extensive
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	detail extensive
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	detail extensive
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	detail extensive

Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	detail extensive
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • Errored blocks—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode. 	detail extensive
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. 	extensive

WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. extensive</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)
Autonegotiation information	<p>Information about link autonegotiation. extensive</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive).

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: <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels

VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive

Transit statistics	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 <i>Common Output Fields Description</i> .	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 <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	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 62: 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 63 on page 785](#) lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 63: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Current address	Configured MAC address.	detail extensive none

Table 63: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 63: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface.</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the <code>ignore-l3-incompletes</code>. • 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	<p>Output errors on the interface.</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; therefore, for Gigabit Ethernet PICs, this number must always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive

Table 63: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters and queue number	CoS queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive

Table 63: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0. 	extensive
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. 	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 63: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
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. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive

Table 63: show interfaces Output Fields (continued)

Field Name	Field Description	Level of Output
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
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
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output Gigabit Ethernet

show interfaces (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: Sendbcst-pkt-to-re

```



```

Addresses, Flags: Is-Preferred Is-Primary
Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
Protocol inet6, MTU: 1500
Max nh cache: 4, New hold nh limit: 100000, Curr nh cnt: 4, Curr new hold
cnt: 4, NH drop cnt: 0
Flags: Is-Primary
Addresses, Flags: Is-Default Is-Preferred Is-Primary
Destination: 2001:db8:/32, Local: 2001:db8::5
Addresses, Flags: Is-Preferred
Destination: 2001:db8:1::/32, Local: 2001:db8:223:9cff:fe9f:3e78
Protocol multiservice, MTU: Unlimited
Flags: Is-Primary

```

show interfaces (link degrade status)

```

user@host> show interfaces et-3/0/0
Physical interface: et-3/0/0, Enabled, Physical link is Down
Interface index: 157, SNMP ifIndex: 537
Link-level type: Ethernet, MTU: 1514, MRU: 0, Speed: 100Gbps, BPDU Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 54:e0:32:23:9d:38, Hardware address: 54:e0:32:23:9d:38
Last flapped   : 2014-06-18 02:36:38 PDT (02:50:50 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : LINK
Active defects : LINK
PCS statistics
  Bit errors           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 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 61 on page 756](#).

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: 99, 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
  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: 34, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled
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 .....
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
Packet Forwarding Engine configuration:
  Destination slot: 1
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %             bps        %      bytes
  0 best-effort           95      950000000    95         0        low    none
  3 network-control       5       50000000     5         0        low    none

```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags : None
Wavelength : 1550.12 nm, Frequency: 193.40 THz
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72
Last flapped : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

```

```

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              0              0              0

  1 expedited-fo            0              0              0

  2 assured-forw            0              0              0

  3 network-cont
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets    0              0
Unicast packets  0              0
Broadcast packets 0              0
Multicast packets 0              0
CRC/Align errors 0              0
FIFO errors       0              0
MAC control frames 0              0
MAC pause frames  0              0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
Total octets      0              0
Total packets    0              0
Unicast packets  0              0
Broadcast packets 0              0
Multicast packets 0              0
CRC/Align errors 0              0
FIFO errors       0              0
MAC control frames 0              0
MAC pause frames  0              0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
OTN alarms       : None
OTN defects      : None
OTN FEC Mode     : GFEC
OTN Rate         : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
Corrected Errors      0
Corrected Error Ratio ( 0 sec average) 0e-0
OTN FEC alarms:      Seconds      Count  State
FEC Degrade          0           0  OK
FEC Excessive         0           0  OK
OTN OC:              Seconds      Count  State
LOS                   2           1  OK
LOF                   67164       2  Defect Active

```

```

LOM                                67164          71 Defect Active
Wavelength Lock                    0              0 OK
OTN OTU:
AIS                                0              0 OK
BDI                                65919          4814 Defect Active
IAE                                67158          1 Defect Active
TTIM                               7              1 OK
SF                                 67164          2 Defect Active
SD                                 67164          3 Defect Active
TCA-ES                             0              0 OK
TCA-SES                             0              0 OK
TCA-UAS                             80             40 OK
TCA-BBE                             0              0 OK
BIP                                0              0 OK
BBE                                0              0 OK
ES                                 0              0 OK
SES                                0              0 OK
UAS                                587            0 OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
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
0 best-effort                       95      9500000000    95      usec      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: Sendbcst-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255

```

Sample Output SRX Gigabit Ethernet

```

user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
Interface index: 135, SNMP ifIndex: 510
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running Down

```

```

Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags      : None
CoS queues     : 8 supported, 8 maximum usable queues
Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms   : LINK
Active defects  : LINK
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: public
  Protocol inet, MTU: 1500
    Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255

```

show interfaces detail (Gigabit Ethernet)

```

user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 0 0 bps
    Output bytes: 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps
  Egress queues: 8 supported, 4 in use
  Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

```

  Queue number:      Mapped forwarding classes
    0                best-effort
    1                expedited-forwarding
    2                assured-forwarding
    3                network-control
  Active alarms   : LINK
  Active defects  : LINK

```

Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Transit statistics:

Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

Security: Zone: public

Flow Statistics :

Flow Input statistics :

Self packets : 0
ICMP packets : 0
VPN packets : 0
Multicast packets : 0
Bytes permitted by policy : 0
Connections established : 0

Flow Output statistics:

Multicast packets : 0
Bytes permitted by policy : 0

Flow error statistics (Packets dropped due to):

Address spoofing: 0
Authentication failed: 0
Incoming NAT errors: 0
Invalid zone received packet: 0
Multiple user authentications: 0
Multiple incoming NAT: 0
No parent for a gate: 0
No one interested in self packets: 0
No minor session: 0
No more sessions: 0
No NAT gate: 0
No route present: 0
No SA for incoming SPI: 0
No tunnel found: 0
No session for a gate: 0
No zone or NULL zone binding: 0
Policy denied: 0
Security association not active: 0
TCP sequence number out of window: 0
Syn-attack protection: 0
User authentication errors: 0

Protocol inet, MTU: 1500, Generation: 150, Route table: 0

Flags: Sendbroadcast-pkt-to-re

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:

150

show interfaces statistics st0.0 detail

```

user@host> show interfaces statistics st0.0 detail
Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
Traffic statistics:
  Input bytes :      528152756774
  Output bytes :     575950643520
  Input packets:    11481581669
  Output packets:   12520666095
Local statistics:
  Input bytes :      0
  Output bytes :      0
  Input packets:     0
  Output packets:    0
Transit statistics:
  Input bytes :      0          121859888 bps
  Output bytes :     0          128104112 bps
  Input packets:     0          331141 pps
  Output packets:    0          348108 pps
Security: Zone: untrust
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
sap vrrp
Flow Statistics :
Flow Input statistics :
  Self packets :      0
  ICMP packets :      0
  VPN packets :      0
  Multicast packets : 0
  Bytes permitted by policy : 525984295844
  Connections established : 7
Flow Output statistics:
  Multicast packets : 0
  Bytes permitted by policy : 576003290222
Flow error statistics (Packets dropped due to):
  Address spoofing:      0
  Authentication failed: 0
  Incoming NAT errors:   0
  Invalid zone received packet: 0
  Multiple user authentications: 0
  Multiple incoming NAT: 0
  No parent for a gate:  0
  No one interested in self packets: 0
  No minor session:      0
  No more sessions:      0
  No NAT gate:           0
  No route present:      2000280
  No SA for incoming SPI: 0
  No tunnel found:       0
  No session for a gate:  0
  No zone or NULL zone binding 0
  Policy denied:         0
  Security association not active: 0
  TCP sequence number out of window: 0
  Syn-attack protection: 0
  User authentication errors: 0
Protocol inet, MTU: 9192
Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
NH drop cnt: 0
Generation: 155, Route table: 0

```

Flags: Sendbroadcast-pkt-to-re

show interfaces extensive (Gigabit Ethernet)

```

user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
Interface index: 135, SNMP ifIndex: 510, Generation: 138
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:00:5e:00:53:01, Hardware address: 00:00:5e:00:53:01
Last flapped : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort        0                0                0
1 expedited-fo       0                0                0
2 assured-forw       0                0                0
3 network-cont       0                0                0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control

Active alarms : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets      0            0
Unicast packets    0            0
Broadcast packets  0            0
Multicast packets  0            0
CRC/Align errors   0            0
FIFO errors        0            0
MAC control frames 0            0

```

```

MAC pause frames          0          0
Oversized frames          0
Jabber frames             0
Fragment frames           0
VLAN tagged frames        0
Code violations            0
Filter statistics:
  Input packet count       0
  Input packet rejects     0
  Input DA rejects         0
  Input SA rejects         0
  Output packet count      0
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 2, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      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: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
Generation: 150

```

show interfaces terse

```

user@host> show interfaces terse

```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	10.209.4.61/18	
gr-0/0/0	up	up			
ip-0/0/0	up	up			
st0	up	up			
st0.1	up	ready	inet		
ls-0/0/0	up	up			
lt-0/0/0	up	up			
mt-0/0/0	up	up			
pd-0/0/0	up	up			
pe-0/0/0	up	up			
e3-1/0/0	up	up			
t3-2/0/0	up	up			
e1-3/0/0	up	up			
se-4/0/0	up	down			
t1-5/0/0	up	up			
br-6/0/0	up	up			
dc-6/0/0	up	up			
dc-6/0/0.32767	up	up			
bc-6/0/0:1	down	up			
bc-6/0/0:1.0	up	down			
d10	up	up			
d10.0	up	up	inet		
dsc	up	up			
gre	up	up			
ipip	up	up			
lo0	up	up			
lo0.16385	up	up	inet	10.0.0.1 10.0.0.16	--> 0/0 --> 0/0

```

lsi                up    up
mtun               up    up
pimd              up    up
pime              up    up
pp0               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 DSO)

```
user@host> show interfaces controller ce1-1/2/3
```

Controller	Admin	Link
ce1-1/2/3	up	up
ds-1/2/3:1	up	up
ds-1/2/3:2	up	up

show interfaces descriptions

```
user@host> show interfaces descriptions
```

Interface	Admin	Link	Description
so-1/0/0	up	up	M20-3#1
so-2/0/0	up	up	GSR-12#1
ge-3/0/0	up	up	SMB-OSPF_Area300
so-3/3/0	up	up	GSR-13#1
so-3/3/1	up	up	GSR-13#2
ge-4/0/0	up	up	T320-7#1
ge-5/0/0	up	up	T320-7#2
so-7/1/0	up	up	M160-6#1
ge-8/0/0	up	up	T320-7#3
ge-9/0/0	up	up	T320-7#4
so-10/0/0	up	up	M160-6#2
so-13/0/0	up	up	M20-3#2
so-14/0/0	up	up	GSR-12#2
ge-15/0/0	up	up	SMB-OSPF_Area100
ge-15/0/1	up	up	GSR-13#3

show interfaces destination-class all

```
user@host> show interfaces destination-class all
```

```
Logical interface so-4/0/0.0
```

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	0	0
(silver	0)	0)
(0)	0)

```
Logical interface so-0/1/3.0
```

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	0	0
(0)	0)


```

silver                                0                                0
(                                     0) (                             0)

```

show interfaces diagnostics optics

```

user@host> show interfaces diagnostics optics ge-2/0/0
Physical interface: ge-2/0/0
Laser bias current                      : 7.408 mA
Laser output power                      : 0.3500 mW / -4.56 dBm
Module temperature                      : 23 degrees C / 73 degrees F
Module voltage                          : 3.3450 V
Receiver signal average optical power  : 0.0002 mW / -36.99 dBm
Laser bias current high alarm           : Off
Laser bias current low alarm            : Off
Laser bias current high warning         : Off
Laser bias current low warning          : Off
Laser output power high alarm           : Off
Laser output power low alarm            : Off
Laser output power high warning         : Off
Laser output power low warning          : Off
Module temperature high alarm           : Off
Module temperature low alarm            : Off
Module temperature high warning         : Off
Module temperature low warning          : Off
Module voltage high alarm               : Off
Module voltage low alarm                : Off
Module voltage high warning             : Off
Module voltage low warning              : Off
Laser rx power high alarm               : Off
Laser rx power low alarm                : On
Laser rx power high warning             : Off
Laser rx power low warning              : On
Laser bias current high alarm threshold : 17.000 mA
Laser bias current low alarm threshold  : 1.000 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6310 mW / -2.00 dBm
Laser output power low alarm threshold  : 0.0670 mW / -11.74 dBm
Laser output power high warning threshold : 0.6310 mW / -2.00 dBm
Laser output power low warning threshold : 0.0790 mW / -11.02 dBm
Module temperature high alarm threshold : 95 degrees C / 203 degrees F
Module temperature low alarm threshold  : -25 degrees C / -13 degrees F
Module temperature high warning threshold : 90 degrees C / 194 degrees F
Module temperature low warning threshold : -20 degrees C / -4 degrees F
Module voltage high alarm threshold     : 3.900 V
Module voltage low alarm threshold       : 2.700 V
Module voltage high warning threshold   : 3.700 V
Module voltage low warning threshold     : 2.900 V
Laser rx power high alarm threshold     : 1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold       : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold   : 0.7940 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.0158 mW / -18.01 dBm

```

show interfaces far-end-interval coc12-5/2/0

```

user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
05:30-current:
ES-L: 1, SES-L: 1, UAS-L: 0

```

```

05:15-05:30:
  ES-L: 0, SES-L: 0, UAS-L: 0
05:00-05:15:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:45-05:00:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:30-04:45:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:15-04:30:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:00-04:15:
...

```

show interfaces far-end-interval coc1-5/2/1:1

```

user@host> run show interfaces far-end-interval coc1-5/2/1:1
Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
05:30-current:
  ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:15-05:30:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:00-05:15:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:45-05:00:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:30-04:45:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:15-04:30:
  ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:00-04:15:

```

show interfaces filters

```

user@host> show interfaces filters

```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet		
			iso		
ge-5/0/0	up	up			
ge-5/0/0.0	up	up	any		f-any
			inet		f-inet
			multiservice		
gr-0/3/0	up	up			
ip-0/3/0	up	up			
mt-0/3/0	up	up			
pd-0/3/0	up	up			
pe-0/3/0	up	up			
vt-0/3/0	up	up			
at-1/0/0	up	up			
at-1/0/0.0	up	up	inet		
			iso		
at-1/1/0	up	down			
at-1/1/0.0	up	down	inet		
			iso		
....					

show interfaces flow-statistics (Gigabit Ethernet)

```

user@host> show interfaces flow-statistics ge-0/0/1.0

```

```

Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 5161
  Output packets: 83
  Security: Zone: zone2
  Allowed host-inbound traffic : bootp bfd bgp dns dvmrp ldp msdp nhrp ospf
pgm
  pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
https ike
  netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
xnm-ssl
  lsping
  Flow Statistics :
  Flow Input statistics :
    Self packets : 0
    ICMP packets : 0
    VPN packets : 2564
    Bytes permitted by policy : 3478
    Connections established : 1
  Flow Output statistics:
    Multicast packets : 0
    Bytes permitted by policy : 16994
  Flow error statistics (Packets dropped due to):
    Address spoofing: 0
    Authentication failed: 0
    Incoming NAT errors: 0
    Invalid zone received packet: 0
    Multiple user authentications: 0
    Multiple incoming NAT: 0
    No parent for a gate: 0
    No one interested in self packets: 0
    No minor session: 0
    No more sessions: 0
    No NAT gate: 0
    No route present: 0
    No SA for incoming SPI: 0
    No tunnel found: 0
    No session for a gate: 0
    No zone or NULL zone binding 0
    Policy denied: 0
    Security association not active: 0
    TCP sequence number out of window: 0
    Syn-attack protection: 0
    User authentication errors: 0
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255

```

show interfaces interval (Channelized OC12)

```

user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
17:43-current:
  LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
  SEFS: 0, UAS: 0
17:28-17:43:
  LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
  SEFS: 0, UAS: 0
17:13-17:28:

```

```

LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:58-17:13:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:43-16:58:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

show interfaces interval (E3)

```

user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
17:43-current:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
17:28-17:43:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
17:13-17:28:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:58-17:13:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
SEFS: 0, UAS: 0
16:43-16:58:
LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

show interfaces interval (SONET/SDH) (SRX devices)

```

user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
20:02-current:
ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
SES-P: 0, UAS-P: 0
19:47-20:02:
ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
ES-P: 267, SES-P: 267, UAS-P: 267
19:32-19:47:
ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
SES-P: 56, UAS-P: 46
19:17-19:32:
ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
SES-P: 0, UAS-P: 0
19:02-19:17:
.....

```

show interfaces load-balancing (SRX devices)

```

user@host> show interfaces load-balancing
Interface State      Last change  Member count
ams0      Up              1d 00:50     2
ams1      Up              00:00:59     2

```

show interfaces load-balancing detail (SRX devices)

```

user@host>show interfaces load-balancing detail
Load-balancing interfaces detail
Interface      : ams0
State          : Up
Last change    : 1d 00:51
Member count   : 2
Members        :
  Interface    Weight  State
  mams-2/0/0   10     Active
  mams-2/1/0   10     Active

```

show interfaces mac-database (All MAC Addresses on a Port SRX devices)

```

user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2

```

MAC address	Input frames	Input bytes	Output frames	Output bytes
00:00:00:00:00:00	1	56	0	0
00:00:c0:01:01:02	7023810	323095260	0	0
00:00:c0:01:01:03	7023810	323095260	0	0
00:00:c0:01:01:04	7023810	323095260	0	0
00:00:c0:01:01:05	7023810	323095260	0	0
00:00:c0:01:01:06	7023810	323095260	0	0
00:00:c0:01:01:07	7023810	323095260	0	0
00:00:c0:01:01:08	7023809	323095214	0	0
00:00:c0:01:01:09	7023809	323095214	0	0
00:00:c0:01:01:0a	7023809	323095214	0	0
00:00:c0:01:01:0b	7023809	323095214	0	0
00:00:c8:01:01:02	30424784	1399540064	37448598	1722635508
00:00:c8:01:01:03	30424784	1399540064	37448598	1722635508
00:00:c8:01:01:04	30424716	1399536936	37448523	1722632058
00:00:c8:01:01:05	30424789	1399540294	37448598	1722635508
00:00:c8:01:01:06	30424788	1399540248	37448597	1722635462
00:00:c8:01:01:07	30424783	1399540018	37448597	1722635462
00:00:c8:01:01:08	30424783	1399540018	37448596	1722635416
00:00:c8:01:01:09	8836796	406492616	8836795	406492570
00:00:c8:01:01:0a	30424712	1399536752	37448521	1722631966
00:00:c8:01:01:0b	30424715	1399536890	37448523	1722632058

```

Number of MAC addresses : 21

```

show interfaces mac-database (All MAC Addresses on a Service SRX devices)

```

user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2

```

MAC address	Input frames	Input bytes	Output frames	Output bytes
00:00:00:00:00:00	1	56	0	0

00:00:c0:01:01:02	7023810	323095260	0	0
00:00:c0:01:01:03	7023810	323095260	0	0
00:00:c0:01:01:04	7023810	323095260	0	0
00:00:c0:01:01:05	7023810	323095260	0	0
00:00:c0:01:01:06	7023810	323095260	0	0
00:00:c0:01:01:07	7023810	323095260	0	0
00:00:c0:01:01:08	7023809	323095214	0	0
00:00:c0:01:01:09	7023809	323095214	0	0
00:00:c0:01:01:0a	7023809	323095214	0	0
00:00:c0:01:01:0b	7023809	323095214	0	0
00:00:c8:01:01:02	31016568	1426762128	38040381	1749857526
00:00:c8:01:01:03	31016568	1426762128	38040382	1749857572
00:00:c8:01:01:04	31016499	1426758954	38040306	1749854076
00:00:c8:01:01:05	31016573	1426762358	38040381	1749857526
00:00:c8:01:01:06	31016573	1426762358	38040381	1749857526
00:00:c8:01:01:07	31016567	1426762082	38040380	1749857480
00:00:c8:01:01:08	31016567	1426762082	38040379	1749857434
00:00:c8:01:01:09	9428580	433714680	9428580	433714680
00:00:c8:01:01:0a	31016496	1426758816	38040304	1749853984
00:00:c8:01:01:0b	31016498	1426758908	38040307	1749854122

show interfaces mac-database mac-address

```

user@host> show interfaces mac-database xe-0/3/3 mac-address (SRX devices)
00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address: 00:00:c8:01:01:09, Type: Configured,
  Input bytes   : 202324652
  Output bytes  : 202324560
  Input frames  : 4398362
  Output frames : 4398360
Policer statistics:
Policer type    Discarded frames  Discarded bytes
Output aggregate      3992386        183649756

```

show interfaces mc-ae (SRX devices)

```

user@host> show interfaces mc-ae ae0 unit 512
Member Links   : ae0
Local Status   : active
Peer Status    : active
Logical Interface      : ae0.512
Core Facing Interface : Label Ethernet Interface
ICL-PL          : Label Ethernet Interface

```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```

user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues     : 8 supported
  Last flapped   : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
    BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2

```

show interfaces policers (SRX devices)

```

user@host> show interfaces policers
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/0       up    up
ge-0/0/0.0     up    up    inet
               up    up    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
               up    up    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
               up    down iso
               up    down 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
r1sq0     On secondary  00:56:12    1sq-4/0/0    1sq-3/0/0    both up

ae0
ae1
ae2
ae3
ae4

```

show interfaces redundancy detail (SRX devices)

```

user@host> show interfaces redundancy detail
Interface      : r1sq0
State          : On primary
Last change    : 00:45:47
Primary        : 1sq-0/2/0
Secondary      : 1sq-1/2/0
Current status : both up
Mode           : hot-standby

Interface      : r1sq0:0
State          : On primary
Last change    : 00:45:46
Primary        : 1sq-0/2/0:0
Secondary      : 1sq-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
1o0.0          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: 0C48,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
Link flags     : Keepalives
CoS queues     : 8 supported
Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
SONET alarms   : LOL, PLL, LOS
SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P

```

show interfaces source-class all (SRX devices)

```

user@host> show interfaces source-class all
Logical interface so-0/1/0.0

Source class          Packets          Bytes
                    (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:
Total bytes      Receive      Transmit
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 diagnostics optics

Syntax `show interfaces diagnostics optics interface-name`

Release Information Command introduced in Junos OS Release 10.1.

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.

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

Related Documentation

- [Understanding Interfaces on page 3](#)

List of Sample Output [show interfaces diagnostics optics on page 833](#)

Output Fields [Table 64 on page 830](#) lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the general order in which they appear.

Table 64: 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).

Table 64: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
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 .
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 .

Table 64: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
Laser rx power low alarm	Displays whether the receive laser power low alarm is On or Off .
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.
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.

Table 64: show interfaces diagnostics optics Output Fields (continued)

Field Name	Field Description
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

show interfaces flow-statistics

Syntax	show interfaces flow-statistics <i><interface-name></i>
Release Information	Command introduced in Junos OS Release 9.2.
Description	Display interfaces flow statistics.
Options	<p>Interface-name —(Optional) Display flow statistics about the specified interface. Following is a list of typical interface names. Replace <i>pim</i> with the PIM slot and <i>port</i> with the port number. For a complete list, see the “Interface Naming Conventions” on page 8.</p> <ul style="list-style-type: none"> • at-<i>pim</i>/0/<i>port</i>—ATM-over-ADSL or ATM-over-SHDSL interface. • br-<i>pim</i>/0/<i>port</i>—Basic Rate Interface for establishing ISDN connections. • ce1-<i>pim</i>/0/<i>port</i>—Channelized E1 interface. • ct1-<i>pim</i>/0/<i>port</i>—Channelized T1 interface. • dl0—Dialer Interface for initiating ISDN and USB modem connections. • e1-<i>pim</i>/0/<i>port</i>—E1 interface. • e3-<i>pim</i>/0/<i>port</i>—E3 interface. • fe-<i>pim</i>/0/<i>port</i>—Fast Ethernet interface. • ge-<i>pim</i>/0/<i>port</i>—Gigabit Ethernet interface. • se-<i>pim</i>/0/<i>port</i>—Serial interface. • t1-<i>pim</i>/0/<i>port</i>—T1 (also called DS1) interface. • t3-<i>pim</i>/0/<i>port</i>—T3 (also called DS3) interface. • wx-slot/0/0—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Traffic Processing on Security Devices • Understanding Interfaces on page 3
List of Sample Output	show interfaces flow-statistics (Gigabit Ethernet) on page 838
Output Fields	Table 65 on page 836 lists the output fields for the show interfaces flow-statistics command. Output fields are listed in the approximate order in which they appear.

Table 65: 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 66 on page 836 .

Table 66: 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.

Table 66: Flow Error Statistics (Packet Drop Statistics for the Flow Module) (continued)

User authentication errors	<p>Packet was dropped because policy requires authentication; however:</p> <ul style="list-style-type: none"> • 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	<p>This counter is incremented for one of the following reasons:</p> <ul style="list-style-type: none"> • 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	<p>The packet was dropped because a valid route was not available to forward the packet.</p> <p>For new sessions, the counter is incremented for one of the following reasons:</p> <ul style="list-style-type: none"> • 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. <p>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:</p> <ul style="list-style-type: none"> • 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	<p>The error counter is incremented for one of the following reasons:</p> <ul style="list-style-type: none"> • 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.

Table 66: Flow Error Statistics (Packet Drop Statistics for the Flow Module) (continued)

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
```

show interfaces queue

Syntax	<pre>show interfaces queue <both-ingress-egress> <egress> <forwarding-class forwarding-class> <ingress> <interface-name interface-name> <l2-statistics></pre>
Release Information	Command introduced in Junos OS Release 15.1X49-D30 for vSRX.
Description	Display class-of-service (CoS) queue information for physical interfaces.
Options	<p>none—Show detailed CoS queue statistics for all physical interfaces.</p> <p>both-ingress-egress—Display both ingress and egress queue statistics.</p> <p>egress—Display egress queue statistics.</p> <p>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.</p> <p>ingress—Display ingress queue statistics.</p> <p>interface-name interface-name—(Optional) Show detailed CoS queue statistics for the specified interface.</p> <p>l2-statistics—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Understanding Class of Service
List of Sample Output	show interfaces queue (vSRX) on page 842
Output Fields	<p>Table 67 on page 840 lists the output fields for the show interfaces queue command. Output fields are listed in the approximate order in which they appear.</p>

Table 67: 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 <i>Common Output Fields Description</i> .

Table 67: *show interfaces queue Output Fields (continued)*

Field Name	Field 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.
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. <ul style="list-style-type: none"> • Low, non-TCP—Number of low-loss priority, non-TCP bytes dropped because of RED. • Low, TCP—Number of low-loss priority, TCP bytes dropped because of RED. • High, non-TCP—Number of high-loss priority, non-TCP bytes dropped because of RED. • High, TCP—Number of high-loss priority, TCP bytes dropped because of RED.
Queue Buffer Usage:	<ul style="list-style-type: none"> • 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.

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: 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 statistics (View)

Syntax	<code>show interfaces statistics <i>interface-name</i></code>
Release Information	Command introduced in Junos OS Release 10.1.
Description	Displays the interface input and output statistics for physical and logical interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Interfaces on page 3
List of Sample Output	show interfaces statistics on page 843

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

```

show interfaces terse zone

Syntax show interfaces terse zone

Release Information Command introduced in Junos OS Release 12.3X48-D20.

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
```

show ipv6 neighbors

Syntax	show ipv6 neighbors
Release Information	Command introduced in Junos OS Release 12.1X45-D10.
Description	Display information about the IPv6 neighbor cache.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ipv6 neighbors on page 689
List of Sample Output	show ipv6 neighbors on page 845
Output Fields	Table 68 on page 845 lists the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear.

Table 68: 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.
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
----------	-------------------	-------	------	-----	----	------------

show lacp interfaces (View)

Syntax `show lacp interfaces interface-name`

Release Information Command modified in Junos OS Release 10.2.

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—***aenumber***
- Redundant Ethernet—***rethnumber***
- 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

Related Documentation

- [Verifying LACP on Redundant Ethernet Interfaces on page 306](#)

List of Sample Output

- [show lacp interfaces \(Aggregated Ethernet\) on page 849](#)
- [show lacp interfaces \(Redundant Ethernet\) on page 850](#)
- [show lacp interfaces \(Gigabit Ethernet\) on page 850](#)

Output Fields [Table 69 on page 847](#) lists the output fields for the `show lacp interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 69: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.

Table 69: show lacp interfaces Output Fields (continued)

Field Name	Field Description
LACP State	<p>LACP state information for each aggregated interface:</p> <ul style="list-style-type: none"> • Role—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> • Actor—Local device participating in LACP negotiation. • Partner—Remote device participating in LACP negotiation. • Exp—Expired state. Yes indicates the actor or partner is in an expired state. No indicates the actor or partner is not in an expired state. • Def—Default. Yes indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. No indicates the operational partner information in use has been received in 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 69: show lacp interfaces Output Fields (continued)

Field Name	Field Description
LACP Protocol	<p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> Link state (active or standby) indicated in parentheses next to the interface when link protection is configured. Receive State—One of the following values: <ul style="list-style-type: none"> Current—The state machine receives 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: <ul style="list-style-type: none"> Fast Periodic—Periodic transmissions are enabled at a fast transmission rate. No Periodic—Periodic transmissions are disabled. Periodic Timer—Transitory state entered when the periodic timer expires. Slow Periodic—Periodic transmissions are enabled at a slow transmission rate. Mux State—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> Attached—Multiplexer state machine initiates the process of attaching the port to the selected aggregator. Collecting 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

show lacp statistics interfaces (View)

Syntax `show lacp statistics interfaces interface-name`

Release Information Command modified in Release 10.2 of Junos OS.
Command introduced in Release 11.1 of Junos OS for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.

Options *interface-name*—(Optional) Name of an interface.

Required Privilege Level view

Related Documentation

- [Verifying LACP on Redundant Ethernet Interfaces on page 306](#)
- [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](#)

List of Sample Output [show lacp statistics interfaces on page 852](#)

Output Fields [Table 70 on page 851](#) lists the output fields for the `show lacp statistics interfaces` command. Output fields are listed in the approximate order in which they appear.

Table 70: show lacp statistics interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated interface value.
LACP Statistics	<p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> • LACP Rx—LACP received counter that increments for each normal hello. • LACP Tx—Number of LACP transmit packet errors logged. • Unknown Rx—Number of unrecognized packet errors logged. • Illegal Rx—Number of invalid packets received.

Sample Output

show lacp statistics interfaces

```
user@host> show lacp statistics interfaces ae0
```

```
Aggregated interface: ae0
```

LACP Statistics:	LACP Rx	LACP Tx	Unknown Rx	Illegal Rx
ge-2/0/0	1352	2035	0	0
ge-2/0/1	1352	2056	0	0
ge-2/2/0	1352	2045	0	0
ge-2/2/1	1352	2043	0	0

show modem wireless firmware

Syntax	<code>show modem wireless firmware <i>interface-name</i></code>
Release Information	Command introduced in Junos OS 15.1X49-D100
Description	Display modem firmware details for the LTE Mini-PIM.
Options	<ul style="list-style-type: none"> <i>interface-name</i>—The LTE interface is <code>cl-x/0/0</code>, where <i>x</i> is the slot number in which the LTE Mini-PIM is installed.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show modem wireless network on page 856
List of Sample Output	show modem wireless firmware on page 854
Output Fields	Table 71 on page 853 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 71: 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 71: show modem wireless firmware Output Fields (continued)

Field Name	Description
Status of SIM	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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. 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
```

show modem wireless network

Syntax	<code>show modem wireless network <i>interface-name</i></code>
Release Information	Command introduced in Junos OS Release 15.1X49-D100.
Description	Display the status of the modem and the status of the network connection for the LTE Mini-PIM.
Options	<ul style="list-style-type: none"> <i>interface-name</i>—The LTE interface is <code>cl-x/0/0</code>, where <i>x</i> is the slot number in which the LTE Mini-PIM is installed.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show modem wireless profiles on page 859 show modem wireless firmware on page 853
List of Sample Output	show modem wireless network on page 857
Output Fields	Table 72 on page 856 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 72: 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: <ul style="list-style-type: none"> Disconnected Calling Connected
Current Service Status	Status of the network connection. The status can be one of the following states: <ul style="list-style-type: none"> 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 72: show modem wireless network Output Fields (continued)

Field Name	Field Description
Current Service Type	One of the following: <ul style="list-style-type: none"> • Circuit switched (CS) • Packet switched (PS) • Combo (CS, PS) • Invalid
Current Service Mode	One of the following: <ul style="list-style-type: none"> • 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

show modem wireless profiles

Syntax	<code>show modem wireless profiles <i>interface-name</i> slot <i>slot-number</i></code>
Release Information	Command introduced in Junos OS Release 15.1X49-D100.
Description	Display the profiles configured on the LTE Mini-PIM.
Options	<ul style="list-style-type: none"> • <i>interface-name</i>—The LTE interface is <code>cl-x/0/0</code>, where <i>x</i> is the slot number in which the LTE Mini-PIM is installed. • <i>slot-number</i>—The slot in which the SIM card is inserted. The value can be either 1 or 2.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show modem wireless firmware on page 853 • show modem wireless network on page 856
List of Sample Output	show modem wireless profiles on page 859
Output Fields	Table 72 on page 856 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 73: 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	<ul style="list-style-type: none"> • 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
```

show oam ethernet link-fault-management

Syntax	show oam ethernet link-fault-management <brief detail> <interface-name>
Release Information	Statement for SRX Series devices introduced in Junos OS Release 9.5.
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
Related Documentation	<ul style="list-style-type: none"> • clear oam ethernet connectivity-fault-management path-database on page 685 • clear oam ethernet connectivity-fault-management statistics • Understanding Ethernet OAM Link Fault Management for SRX Series Services Gateways on page 355 • Example: Configuring Ethernet OAM Link Fault Management on a Security Device on page 357
List of Sample Output	show oam ethernet link-fault-management brief on page 865 show oam ethernet link-fault-management detail on page 865
Output Fields	Table 74 on page 861 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 74: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Status of the established link. <ul style="list-style-type: none"> • Fail—A link fault condition exists. • Running—A link fault condition does not exist. 	All levels

Table 74: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
Discovery state	State of the discovery mechanism: <ul style="list-style-type: none"> • Passive Wait • Send Any • Send Local Remote • Send Local Remote Ok 	All levels
Peer address	Address of the OAM peer.	All levels
Flags	Information about the interface. <ul style="list-style-type: none"> • 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
Remote entity information	Remote entity information. <ul style="list-style-type: none"> • 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

Table 74: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
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
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
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

Table 74: show oam ethernet link-fault-management Output Fields (continued)

Field Name	Field Description	Level of Output
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
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 74: 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

```

show poe controller (View)

Syntax	show poe controller
Release Information	Command introduced in Junos OS Release 9.5.
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
Related Documentation	<ul style="list-style-type: none"> Example: Configuring PoE on All Interfaces on page 370
Output Fields	Table 75 on page 866 lists the output fields for the show poe controller command. Output fields are listed in the approximate order in which they appear.

Table 75: 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 index	Maximum power	Power consumption	Guard band	Management
0	150.0 W	0.0 W	0 W	Static

show pppoe interfaces

Syntax	show pppoe interfaces <brief detail extensive> <pp0.logical>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display session-specific information about PPPoE interfaces.
Options	<p>none—Display interface information for all PPPoE interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>extensive—(Optional) Display information about the number of packets sent and received and the number of timeouts during a PPPoE session.</p> <p>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.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Ethernet Interfaces on page 259
List of Sample Output	show pppoe interfaces on page 869 show pppoe interfaces brief on page 869 show pppoe interfaces detail on page 869 show pppoe interfaces extensive on page 869
Output Fields	Table 76 on page 867 lists the output fields for the show pppoe interfaces command. Output fields are listed in the approximate order in which they appear.

Table 76: 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).

Table 76: show pppoe interfaces Output Fields (continued)

Field Name	Field Description
Configured AC name	Configured access concentrator name.
Session AC name	Name of the access concentrator.
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 <i>hh:mm:ss</i> .
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.
Packet Type	<p>Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <ul style="list-style-type: none"> • 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	<p>Timeouts that occur during the PPPoE session:</p> <ul style="list-style-type: none"> • 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 76: show pppoe interfaces Output Fields (continued)

Field Name	Field Description
Receive Error Counters	<p>Error counters received during the PPPoE session:</p> <ul style="list-style-type: none"> • 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          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	

show pppoe statistics

Syntax	show pppoe statistics <logical-interface-name>
Release Information	Command is introduced in Junos OS Release 9.5.
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
Related Documentation	<ul style="list-style-type: none"> • show pppoe interfaces on page 867 • Understanding Ethernet Interfaces on page 259
List of Sample Output	show pppoe statistics on page 872
Output Fields	Table 77 on page 871 lists the output fields for the show pppoe statistics command. Output fields are listed in the approximate order in which they appear.

Table 77: show pppoe statistics Output Fields

Field Name	Field Description
Active PPPoE sessions	Total number of active PPPoE sessions.
Packet Type	Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors: <ul style="list-style-type: none"> • 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.

Table 77: show pppoe statistics Output Fields (continued)

Field Name	Field Description
Timeout	<p>Timeouts that occur during the PPPoE session:</p> <ul style="list-style-type: none"> • 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.
Receive Error Counters	<p>Error counters received during the PPPoE session:</p> <ul style="list-style-type: none"> • 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

```


show poe telemetries

Syntax	show poe telemetries <interface <i>interface-name</i> count <i>number</i> > <count <i>number</i> interface <i>interface-name</i> >
Release Information	Command modified in Junos OS Release 12.3X48-D10.
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	<ul style="list-style-type: none"> • Interface <i>interface-name</i>—Display telemetries for the specified PoE interface. • count <i>number</i>—Display the specified number of telemetries records for the specified PoE interface.
Required Privilege Level	View
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring PoE on All Interfaces on page 370
Output Fields	Table 78 on page 873 lists the output fields for the show poe telemetries interface command. Output fields are listed in the approximate order in which they appear.

Table 78: 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.
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
```

S1 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

S1 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

show services accounting

Syntax	<pre> 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 </pre>
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.
Description	Display sampled accounting service.
Options	<ul style="list-style-type: none"> • aggregation—Display aggregation information. • errors —Display error statistics. <ul style="list-style-type: none"> • inline-jflow — Display service accounting inline flow monitoring parameters. • fpc-slot <i>slot number</i>— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring. • flow—Display flow information. <ul style="list-style-type: none"> • inline-jflow — Display service accounting inline flow monitoring parameters. • fpc-slot <i>slot number</i>— 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. <ul style="list-style-type: none"> • inline-jflow — Display service accounting inline flow monitoring parameters. • fpc-slot <i>slot number</i>— Display Flexible PIC Concentrator (FPC) slot for inline flow monitoring. • usage—Display CPU usage.
Required Privilege Level	view

Related Documentation	<ul style="list-style-type: none">• Configuring Flow Aggregation to Use Version 9 Flow Templates on page 32
List of Sample Output	show services accounting status inline-jflow on page 876 show services accounting errors inline-jflow on page 876 show service accounting flow inline-jflow on page 876
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
```

show services accounting aggregation (View)

Syntax show services accounting aggregation

Release Information Command introduced in Junos OS Release 10.4.

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

Related Documentation

- [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 32](#)

show services accounting aggregation template (View)

Syntax show services accounting aggregation template

Release Information Command introduced in Junos OS Release 10.4.

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

Related Documentation • [Configuring Flow Aggregation to Use Version 9 Flow Templates on page 32](#)

show services accounting flow-detail (View)

Syntax	show services accounting flow-detail
Release Information	Command introduced in Junos OS Release 10.4.
Description	Display flow detail
Options	<ul style="list-style-type: none">• 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
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Aggregation to Use Version 9 Flow Templates on page 32